

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

AI Memo No. 344

February 1976

HOW NEAR IS NEAR?

a

near

specialist

Murray Elias Denofsky

Abstract

This paper presents a system for understanding the concept of near and far, weighing such factors as purpose of the judgement, dimensions of the objects, absolute size of the distance, and size of the distance relative to other objects, ranges, and standards. A further section discusses the meaning of phrases such as very near, much nearer than, and as near as. Although we will speak of near as a judgement about physical distance, most of the ideas developed will be applicable to any continuous measurable parameter, such as size or time. An adaptation for rows (discrete spaces) is made as well.

"It's not the pale moon that
excites me, that thrills and
delights me, oh, no -- it's
just the nearness of you."

- Ned Washington
(popularised by Glenn Miller)

TABLE OF CONTENTS

PART I: PRELIMINARIES

Scope and Method	7
Origins and Related Work	8
Summary of the Paper	9
Some Remarks on Use of the Geometric Mean	12
How Vague is <u>Near</u> ?	13
How Transitive is <u>Near</u> ?	16
Linguistic Usage:	22
Symmetric & Asymmetric	22
Pair-definite, Single-definite & Semi-definite	22
Locative & Comparative	22
Examples	24
Determining Usage	26
Strictness and Laxity	27
Selection of Objects	27

PART II: DETERMINING THE NEAR THRESHOLD

Factors Involved in Determining The Threshold - A Motivation:	
Definitions	30
Range	31
Object Size	31
Local-Global	31
Absolute Size	31
Standard	32
Distance between Cars on a Road	32
Continuous-Discrete	32
Continuous Case:	
Range and Standards:	
Purpose & Range, Examples	33
Standard	34
Dimensionality of Range	35
Long Ranges	36
Examples	37
Large and Small Domains	37
Selecting Range and Standard	38
Object Size:	
Limitations of Range & Standard	38
Object Size	39
Poles and Corridors	41
Reconciling Different Thresholds:	
Geometric Mean	49
Local-Global Mean Rule	49
Spatial Threshold Mean Rule	49
Two-Mean Rule	49

Approximation by Arithmetic Mean	49
Relative Size of Thresholds	49
Choosing a Discriminating Threshold	49
Planetary Orbits	51
Default Thresholds	51
Cheap Computation of Geometric Mean	52
Relative Precedence	52
Absolute Size	53
Strict and Lax Thresholds	53
Distance between Cities	54
Furniture and Walls	55
Discrete case:	
Large Number of Objects:	
Range Rule and Four-Object Rule	55
Offices in a Building	57
Pages in a Book	57
Lines on a Page	57
Books on a Bookshelf	58
Performance of Continuous vs. Discrete Rules	58
Long and Short Bookshelves	58
Suburbs of a City	58
Towns in Massachusetts	59
Small Number of Objects:	
How the Mean Rule Does for Small n	59
Three- and Two-Choice Cases	61
Axiomatic Approach	62
A Common Basis for Continuous & Discrete Rules	65
 PART III: FURTHER DEVELOPMENTS AND APPLICATIONS	
Consecutive Nearness Thresholds	67
Comparison Thresholds	69
Dimensions other than Distance:	
Time	69
Animal Size	70
Non-Measurable Spaces:	
Emotional Closeness	72
Grouping by Nearness: A Failure of the Theory	72
 Appendix: Mathematical Capabilities of the Human Perceptual Mechanism	74
 Acknowledgements and References	75

LIST OF FIGURES

(Starred figures display major results)

Figures 1a-b	*The Vagueness Band	14
Figures 2a-c	Alternative Gradual Thresholds	17-19
Figure 3	*Transitivity Table for 1 Dimension	20
Figure 4	Two-Dimensional Transitivity	21
Figure 5	*Dimensions of Usage, with Examples	23
Figure 6	*Range Threshold	34
Figures 7a-c	Two-Dimensional Ranges	35
Figure 8	*Object Size Threshold	40
Figures 9a-n	*Object Size Threshold Examples	42-48
Figure 10	*Summary of Threshold Determination	50
Figure 11	*Comparison of Four Rules for Discrete Case	56
Figure 12	*Small Number of Objects	60
Figure 13	Comparison of 2- and 3-Choice Cases	63
Figure 14	*Consecutive Nearness Thresholds	68
Figure 15	*Comparison Thresholds	70
Figures 16a-b	Grouping by Near	73

PART I: PRELIMINARIES

SCOPE AND METHOD

scope

The question of near falls roughly into two parts. One part is how to put together various kinds of geometric information to decide the near threshold. The other is how to extract the above information from context and general knowledge. This paper will deal chiefly with the first part. It is therefore in the tradition of some recent high-level vision work, and will deal more with geometry and thresholds than with language understanding or structure of knowledge, though there is a section on different linguistic usages of near.

summary

The thresholding system defines several geometric parameters of a situation: range, standard distance, and object size. The first two of these may have both local (specific to the present situation) and global (typical) values, and the object size information takes some account of shape as well. From each of these, a near threshold may be defined. These different thresholds are then "linearly" reconciled by taking their geometric means.

scope

While I do not specify in detail how to make the appropriate choices for each of these parameters in a situation, I do show thru many examples that commonsense choices lead to a good threshold. Thus the system provides a framework to guide selection of relevant geometric information from a situation by telling what kinds of parameters are desired, and then it tells what to do with these parameters.

example

Let's consider an example. If a motorist asks me whether MIT is far from Harvard, I can use the fact that people driving by car in a city have a range (expected maximum distance) of the diameter of the city, or about 16 miles for the Boston area, to infer that near means one-eighth of this, or 2 miles. Thus Harvard is near MIT to a Boston motorist, barely. If a friend of mine in another city asked me the same question out of curiosity, so that the method of transportation was unclear, I could again use the range to get a threshold. Since he knows only that the two places are both in the Boston area, the range would again be the diameter of the Boston area, and the answer would once more be 2 miles. This shows how a salient geometrical parameter can often be successfully used as a "default" range, regardless of the exact purpose of the near judgment. On the other hand, if a pedestrian asked me the question, or if my friend already knew both

places were in Cambridge, I would use smaller ranges appropriate to these situations.

myself
the
only
subject

Although I give many examples in this paper to illustrate the theory, I have not conducted empirical studies of the meaning of near on any human subjects other than myself. However if such studies reveal that the theory does not quite reflect general usage, I have at least provided a basic vocabulary of concepts which, with minor adjustments, is likely to be able to represent the general usage.

ORIGINS AND RELATED WORK

Minsky

In <Minsky 1974> attention is called to the lack of usefulness of continuous-range numerical data. In using a measurement, what is ultimately used, very often, is not the exact value, but some qualitative judgement based on this value. This is represented in English by such words as "near," "far," "very near," "small," "big," etc.

Winston

Methods for dealing with this problem are suggested by ideas in <Winston 1970> and <Freiling 1973>. Winston develops a system of decision rules based on fixed numerical cutoffs or thresholds, for deciding when a set of objects in a scene are similar enough to each other, close enough together, or sufficiently similarly oriented, to be considered a "group." Freiling applies the same principles to defining such concepts as "above," "in front of," "standing in the way of," and "surrounded by."

Freiling

Kahn

Two more recent papers developing these methods are the master's theses of Hollerbach and Kahn. Hollerbach applies thresholds to height-width and neck-body ratios of vases to decide such traits as "narrow" and "shallow," which he uses to classify the vase by type. He also uses five discrete categories for curvature and four for rate of change of curvature.

Hollerbach

thresholds
of
 2^n

Kahn, in his time-specialist, represents such terms as "several" and "many" by numbers with "fuzz," or expected errors. Freiling uses one-fifth as a near threshold, and Kahn defines "nearly" or "almost" as one-ninth. My range threshold of one-eighth falls in between theirs. This value has the advantage of being a power of 2, fitting into my system of consecutive nearness thresholds, each double the next (quite near, very near, etc.). Hollerbach uses 2 for about half of his critical ratios, and I use it also in deciding whether a dimension of an object is long or short, in comparison to the longest dimension.

Kahn does not attempt to define the terms "recently" and "a

solving
Kahn's
problem

while ago," pointing out that these terms are too context-dependent to be handled by his methods. I show in DIMENSIONS OTHER THAN DISTANCE how my theory can be adapted very directly to yield a satisfactory treatment of these phrases, which mean essentially "near the present time."

SUMMARY OF THE PAPER

GM
of
several
thresholds

Several different near thresholds are defined, each based on a different kind of evidence. When more than one kind of evidence is available, and no one of them is known to be most relevant, the thresholds are combined by taking their geometric mean (GM), two at a time. GM is used instead of AM (arithmetic mean) because AM is insensitive to the order of magnitude of the smaller quantity being averaged, thus throwing away most of the information from it. A cheap method of computing approximate GM is illustrated.

a
hierarchy
of
consecutive
thresholds

Far is defined as four times near, and other small integral powers of 2 are used to define quite near, very near, not quite near, quite far, etc., giving a hierarchy of consecutive nearness thresholds. A related approach is used to define a system of comparative nearness thresholds, e.g., nearer than, as near as, nearly as far as, etc. The theory is shown to be applicable to time and animal size, as examples of domains other than physical distance.

examples

Copious examples are worked out to show that the theory gives reasonable results in a wide range of applications drawn from everyday life (usually within $\pm 25\%$ of where a human would place the threshold, I claim). Small-domain examples include nearness of lines on a page of print to one another or to the top of the page, and nearness of pages in a book. Larger examples include books on bookshelves of various sizes, nearness of a platform standee to the tracks at an MBTA station, nearness to a pole, nearness of furniture in a room, and nearness to the walls of a long corridor. Still larger examples include nearness of offices on a floor of a building, of buildings on the MIT campus, of an apartment to the MBTA, of cars on a road, of Boston suburbs to one another, of Massachusetts towns, New England and American cities, and planetary orbits. The geographical and astronomical categories illustrate the utility of the GM rule in domains where there is a great difference between object size and object distance.

discrete
vs.
continuous

Different thresholds are defined to handle cases where object size and distance vary continuously, such as nearness of furniture in a room, and "discrete" cases, where objects are closely-packed and/or evenly-spaced, such as books in a bookcase or telephone poles along a road. In the discrete case, distance can be expressed in terms of how

continuous
case:

many objects apart two objects are, and one can tell if any objects are missing from the scene. In the continuous case, distance is measured between the nearest points of the two objects, while in the discrete case, it is measured from the centers.

range
threshold

In continuous case, three kinds of threshold are used: range, object size, and standard. Range threshold is one-eighth of the way from the minimum possible distance to the maximum possible distance. Special attention is given to two-dimensional ranges, both square and elongated.

object
size
threshold

A crude object size threshold can be defined as equal to the diameter of the larger object. I present a more refined version which takes into account the sizes of all six dimensions that the two objects possess between them, but which avoids computing products or square roots. A 3 x 3 table gives factors by which the largest dimension of the larger object should be multiplied to get the threshold. The factors are integral powers of 2 between 1/4 and 4, depending on which of three absolute size categories the largest dimension is in, and on how many of the six dimensions are "large," i.e., greater than half the largest dimension.

standard
threshold

Standard threshold is based on the distance between the members of pairs of "adjacent" objects similar to the pair under discussion and in the same physical neighborhood. It equals this distance or half of it, depending on whether the emphasis is on typical nearness or unusual nearness. The latter would be indicated by such phrases as "so near" or "too near."

local
&
global

For range and standard thresholds, both a local and a global threshold may exist. A local threshold is based on the data of the particular example being discussed, while a global threshold is based on typical values of data in this kind of example, as retrieved from general knowledge. If both a local and a global standard exist, we take their GM to get the standard threshold, and similarly for range. For example, in discussing how closely the car behind is following us on the road, our threshold will be affected by the nearness of other cars to each other in our neighborhood on the road (local standard), since we will have a stricter threshold in a traffic jam. But it will also be affected by the usual nearness of cars on a road of this kind (global standard).

geometric
mean

(GM)

If range, object size, and standard threshold all exist, we take the GM of the first two to get the "spatial threshold," and then the GM of that with the standard threshold. Standard threshold is given more weight because it is a more direct kind of evidence, not requiring computation of complicated spatial considerations that may or may not be relevant. A threshold should be omitted from the mean if its inclusion yields a mean threshold outside, or almost outside, the range.

It's
not so
bad.

Although these rules seem complicated and riddled with ad hoc corrections, the computation is actually quite simple for most examples, since only a few of the rules will be applicable. In particular, it is often unnecessary to take means, since only one threshold may be available, or all the thresholds may be about equal. If they differ by only a factor of about 2, the AM may be harmlessly substituted for the GM.

Discrete cases:

4-object rule
mean rule

In the discrete case, I use the four-object rule, which states simply that two objects in a row of objects are near (or one object is near the end of the row) if they are four objects apart (i.e., three objects in between them). This is adequate when the size of the row is moderate, about two or three dozen objects in length. For the general case, I take the GM of 4 and the range threshold, which is $1/8$ of the total number of objects, giving a threshold of $\sqrt{n/2}$.

2-choice
vs.
3-choice

A discussion is given of two-choice cases (near one end or the other) and three-choice cases (near an end versus near the middle), and an axiomatic treatment of these for small number of objects is given, based on commonsense properties of near. It is shown to be very close to the square root rules.

usage

I define several dimensions of usage of near, including symmetry, definiteness, and locative-comparative. These usages are likened to Martin's distinction between different senses of prepositions such as "with." Clues for identifying usage are discussed, and the relevance of usage to answering questions about near is illustrated.

Vagueness

The paper also includes discussions of vagueness and transitivity. Three kinds of vagueness are distinguished; ambiguity of meaning, useful conceptual vagueness (Kahn's "fuzz"), and statistical vagueness due to random variation between observations. The last may be reducible by repeated observations, but not the first two.

transitivity

From simple statistical considerations, amended by a psychological observation about the suggestive effect of transitivity situations, a realistic theory of transitivity is developed. Several more complicated approaches are then compared, and it is shown that nothing is gained by using normal distributions instead of uniform distributions in defining transitivity. An example is given of how to apply the theory to compute transivities in two-dimensional ranges, as well as a table of one-dimensional transivities for sequences of objects both in order of increasing distance and in random order.

SOME REMARKS ON USE OF THE GEOMETRIC MEAN

My use of the geometric mean to reconcile thresholds derived from different sorts of evidence, when these are equally relevant (as I show is a good default assumption when more than one kind of evidence is present) may at first sight be criticized as "too mathematical," too "simple-minded," or "not what people probably do." I would like to say a number of things in answer to these objections.

criticisms

Firstly, the method is not a sophisticated mathematical concept, but simply linear interpolation, the simplest way of combining two numbers. The GM is a more natural form of interpolation here than the arithmetic mean because the latter is insensitive to the order of magnitude of the smaller quantity; changing this quantity by a factor of 1/10 or 1/100 may have no appreciable effect on the AM, but does affect the near threshold. The GM is always more relevant in situations where ratios rather than differences are what count, and where there is a natural zero to act as a reference point (zero distance, in our case).

GM
better
than
AM

Computationally the GM is not really much more complex than the AM. Since near thresholds are computed only to within $\pm 25\%$, the amount of computation required is small in either case (see page 52).

computational
complexity

These observations are probably true for people in particular, since we are more sensitive to ratios than differences; the Weber-Fechner Law (see page 74) says we respond to the logarithm of a stimulus, i.e., equal ratios have equal effect. This law provides a rationale for my "exponential" scale of consecutive nearness thresholds, (near, quite near, very near, etc.), in which each threshold is half the previous, and is the GM of the thresholds on each side of it. It is likely that while higher level thought processes (including many higher level vision processes) are symbolic, metrical operations, or some operation capable of computing GM, plays a role too. Since this can be done by taking logarithms and bisecting the resulting interval, and we have just seen that the brain takes logarithms all the time, such a process would be a lot easier to envision than, say, a process for analog rotation of cubes in the mind.

evidence
that
people
can
do it

The simple linear idea works in reconciling thresholds because the more difficult considerations have already been taken care of in computing the individual thresholds. At the present stage, the thresholds are numbers and must be combined to yield a result that depends continually on both. Only when this final mean threshold is applied, is the number converted to a verbal judgment.

linearity

Minsky's
frames

My use of interpolation to reconcile several ideas based on more sophisticated knowledge is similar in spirit to Minsky's frame system for perceiving rooms and cubes. He suggests linear interpolation between several typical, but qualitatively different, views. It is important to distinguish this use of interpolation as a final step in dealing with data which has been sufficiently cogitated over, from, say, Greenblatt's static evaluator for chess, where linear combination is used as a first approximation in a situation that requires further thought.

elegance
&
versatility

My examples show that the GM of range, object size and standard thresholds is an elegant representation of the meaning of near over wide quantitative and qualitative ranges, including distance between cities, between furniture in a room, and between pages in a book. Whether people indeed compute GM in either a digital or analog fashion, or whether they memorize a table of values or a large collection of frames, is another question. It seems unlikely they could have enough frames to interpolate smoothly along several dimensions simultaneously, without some form of linear smoothing at least between frames, if not as the basic formula for the whole process. And such a system of frames with smoothing, especially since it would probably involve separate frame systems for different contexts, would seem like an unwieldy way of doing a problem which I have shown can be handled in all contexts so simply. Nevertheless it is certainly possible that people use such a patchwork method, in view of its applicability to a wide range of other problems.

HOW VAGUE IS NEAR?

a
vagueness
band

In deciding for which, of various possible distances between two objects, the two objects would be considered near, there is an interval or band in which the distance could be considered either near or not near. This band separates the two regions in which the distance is clearly near or clearly not near. In estimating the near threshold, a person will choose some point in this band.

estimating
it

I estimate the width of the band, for situations in which the range of possible distances is not inordinately large, to be about 50% of the mean threshold T (the center of the band), extending about 25% on each side of this mean (see Figure 1a). Thus the lowest acceptable value of the threshold is $3/4$ of the mean threshold, and is $3/4 \times 4/5 = 3/5$ of the highest acceptable value. I think it is reasonable that the lowest and highest acceptable values should differ by about a factor of 2.

These ratios of $3/4$ and $4/5$ are close to $1/\sqrt{2}$ ($\sim .71$). This suggests that, although the $\pm 25\%$ vagueness band is convenient for




Fig. (1a) - Vagueness band representing the $\pm 25\%$ statistical vagueness in the near threshold.

practical purposes, a more natural definition would be to set the band limits at $1/\sqrt{2} T$ and $\sqrt{2} T$, so that the ratio of extreme acceptable thresholds is exactly 2, consistent with our hierarchy of consecutive nearness thresholds differing by ratios of 2. The highest acceptable value of each threshold is then the lowest acceptable value of the next higher threshold (Figure 1b; for further remarks on this, see "a ratio of 2" in CONSECUTIVE NEARNESS THRESHOLDS). Also, by using ratios instead of differences, the error allowed on the far side is greater, which makes sense, since there is more room for error there. $\sqrt{2}$ (≈ 1.414) means we are allowing as much as a 41% error on the far side, which may be a little excessive, but this extreme value will rarely occur, since errors are distributed normally.

rationale

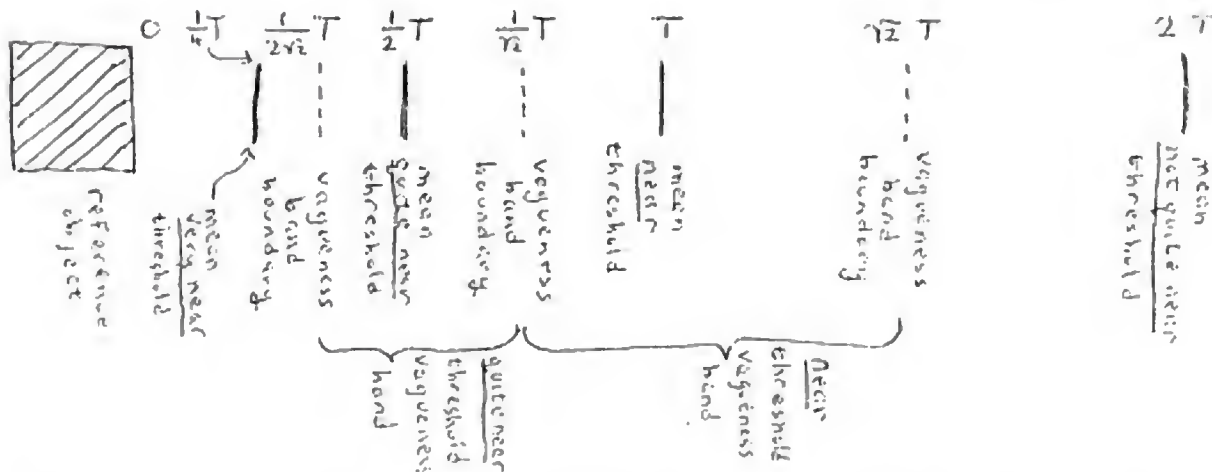


Fig. (1b) Consecutive Nearness Thresholds and Vagueness Bands

Each threshold has an uncertainty in location extending out to the geometric means of its mean position and the mean positions of the neighboring thresholds. See CONSECUTIVE NEARNESS THRESHOLD for related material (page 67).

ambiguity

It would be interesting to test the above ideas with human subjects. I think that if an empirical study shows a band much wider than 25%, it is likely that the subjects are attributing different meanings to near. To avoid this ambiguity, one must make the context or purpose of the judgment sufficiently clear.

wider
bands

But when the possible distances vary over a great range (at least 2 or 3 orders of magnitude), one may wish to increase the bandwidth significantly over 50% of the threshold. Thus even if one has some idea of purpose (e.g., to search for planets), it is, quite likely very vague to say "The spaceship was near Algal." One person may set the threshold at five light-years, another at one-tenth that. In general, we can set the bandwidth equal to the ratio of consecutive nearness thresholds, as defined on page 67.

meaning
of
the
band

There are two ways in which we can view the 25% vagueness band. It may be due to an innate uncertainty in the threshold; there is no good reason why the threshold should be located at one point rather than a nearby point. Or it may be that there exists an ideal threshold which could be approached as closely as we wish by taking the mean of a large number of human judgments, provided that we can avoid the ambiguity problem. The vagueness in the second case would be a statistical error (statistical vagueness). (Since we decided that human judgments almost always fall in this band, the band would correspond to at least two standard deviations (95%) of the error distribution, which would be some variant of a normal distribution.) While it is difficult to disentangle these different kinds of vagueness, it is convenient for us to assume an ideal threshold, since it is easier to design a theory which predicts one than to design a theory which predicts a vague threshold. We can always add the $\pm 25\%$ vagueness caveat afterwards.

statistical
vagueness

how well
our
system
does

While we do not necessarily expect our theory to really predict the center of the band, we do hope that it will at least fall within the band. In most of the examples I have tried, the theory does do this well. In the remaining examples, it comes close. In the few cases when it falls outside the band, one can usually spot some logical consideration unique to the situation, which when added to the general theory gives a better result (e.g. the corridor example on page 39). The reader may judge for himself how justifiable these "ad hoc" considerations are. The real world is bound to include occasional special circumstances, especially from the point of view of today's ignorant computers, and so I do not consider taking account of them in this way to be "fudging," as long as I can show a good rationale which clearly is not applicable elsewhere.

how
accurately
we should
compute

This value of $\pm 25\%$ also guides us in deciding how accurately we should compute the threshold in our theory. Inaccuracies in computing the theoretical prediction can stem both from failure to

interpolate smoothly the theory, which sets up object size and shape categories based on integral powers of two, and from arithmetical approximations in the computation. There is no point in computing the threshold to much better than $\pm 25\%$.

conceptual
vagueness

When I call an object near, I am saying it is within the near threshold. I may also be implying sometimes that it is not within the next smaller threshold, quite near. In any case, there is still some uncertainty in the object's location. Unlike ambiguity and statistical vagueness, this uncertainty is a desirable characteristic of the near concept, giving it its usefulness. We can call it conceptual vagueness. It is equivalent to Kahn's "fuzz."

Summary:
(3 kinds
of
vagueness)

Let us review the three kinds of vagueness defined above. Conceptual vagueness or "fuzz" is an uncertainty deliberately incorporated into the definition of a concept. Statistical vagueness is a generally smaller uncertainty caused by human inconsistency, and reducible by averaging repeated judgments when a more precise value would be meaningful. Ambiguity is a larger uncertainty caused by lack of knowledge of the purpose of the judgment, or, if you like, by the fact that one word represents a number of different concepts.

HOW TRANSITIVE IS NEAR?

Another problem that suggests itself in defining near is how to ensure the right amount of transitivity. For example, if we know that A is near B, and B is near C, then we would expect that A is probably near C, but clearly we would not want such a chain of reasoning to be extended indefinitely.

uniform
distribution

Most definitions of near will be based on some cutoff point. We might define, for example, two objects to be near each other whenever their distance is less than some threshold T. If we assume that (NEAR A B) implies that the position is distributed uniformly within the interval (0, $\pm T$) then (NEAR A B) and (NEAR B C) imply that there is a 50% chance of (NEAR A C), if the three objects are in alphabetical order on a straight line. (The sum of the two means, each T/2, gives the mean sum, T. Also the sum of two random variables with symmetric distributions is itself symmetric. Therefore its mean is its median. So the chance is 50% that the sum is less than T.)

alphabetical
order
on a line

random order
on a line

If they are in random order on a straight line, we get 75%. (There is equal chance that BC must be added to, or subtracted from AB. If added we have the alphabetical case, getting $1/2 \times 50\% = 25\%$, and if subtracted, we get $1/2 \times 100\% = 50\%$, giving a total chance of 75% for (NEAR A C).)

not good
enough

Intuitively, it would seem that, if we have (NEAR A B) and

(NEAR B C), the chances of (NEAR A C) are a lot closer to 100% than is given by the above results. There are two ways in which we can try to modify our theory to account for this.

first
solution:
gradual
threshold

The first method that suggests itself is to assume that (NEAR A B) does not imply a uniform distribution, since the threshold is vague, indicating a gradual cutoff. Thus the mean distance is smaller than in the uniform case, and the sum of two is more likely to be near. The curve will be something like the one in Figure 2a.




Fig. (2a) - Vague threshold fails to explain the laxity of transitive nearness.

This curve gives the chance that a given distance will be called near. The reverse problem, the probability that a distance has a given value, given that someone has told us it is near, has the same curve, but normalized to have unit area. This follows from Bayes' Law.

it backfires

But this curve has a higher mean than did the uniform distribution, because it has a tail outside the threshold as large as the missing area inside the threshold!

a
better
idea

The second idea is to assume that when people judge the nearness of AC, they use a laxer threshold, since they "expect" AC to be near. If one is so ignorant that one has to depend on transitivity to locate an object, then even a lax threshold is a help, and too strict a threshold would be a hindrance to the speaker in selecting a

transitive
"suggestibility"

suitable reference object. Thus we shall assume that for AC (or AD, etc.) to be near, they need only be at a distance less than twice the near threshold. This is the not quite near or almost near threshold (see CONSECUTIVE NEARNESS THRESHOLDS). Thus, in the alphabetical case, AC has a 100% chance of almost near, and AD an 83% chance (5/6), and in the random order case, AC is 100% again, and AD 96% (23/24) (computed from the sum of uniform distributions).

prelude

The above analysis is adequate, but my curiosity led me to explore some further statistical approaches to this problem. These occupy most of the rest of this section.




Fig. (2b) - A closer look at gradual threshold built by "compound" distribution. As a way of incorporating statistical vagueness, it does not change the results much from the uniform distribution.

"compound"
distribution

If we also want to take into account the vagueness of near, we could try using both of the above ideas. Since we want the threshold to be spread over a band of width $.5 T$ (see HOW VAGUE IS NEAR?), with very little chance of being anywhere else, we can assume the threshold is normally distributed around the mean threshold, with the band covering two standard deviations on each side of the mean (95% of the normal curve area). For distances less than those in the band, the distribution is uniform, giving us a "compound distribution." Figure 2b shows that, if AB is near, then there is about a 1/16 chance that

not
worth
it

it is beyond the mean threshold. The chance that it will not be judged near in a second, independent judgment is then also roughly 1/16, since the new threshold will on the average be somewhere near the mean threshold. Thus, if A is near B, there is about a 6% chance that it isn't! But the chance that it is almost near is virtually 100%. For the random order case, the chance that AC is near is probably about 73%, and almost near, 99%. (I estimated these figures by summing the components of the compound distribution.) These figures are almost the same as for the uniform distribution.

The higher order transitivities can be computed from the distribution of the sum of n independent random variables. The computation is non-trivial, but it only needs to be done once.

advantages
of a
pure normal
distribution

An ingenious possibility is that the computation could be made a lot easier if we could use, instead of the uniform or compound distributions, a pure normal distribution, since the sum of normal distributions is also normal. This can be done by assuming that the position of B is distributed normally about A, with a standard deviation about equal to the mean threshold, as in Figure 2c. The probabilities are as shown in the bottom row of the table in Figure 3, which summarizes our results so far.




Fig. (2c) - An ingenious way of using a pure normal distribution to represent near. It allows easy computation of transitivity, but dictates an inflexible, large vagueness.

disadvantages The normal values are much lower than the others, and cannot be taken as an approximation to them, unfortunately. They correspond

TRANSITIVITY TABLE FOR 1 DIMENSION

case	method	<u>AB</u>		<u>AC</u>		<u>AD</u>	
		<u>near</u>	<u>almost near</u>	<u>near</u>	<u>almost near</u>	<u>near</u>	<u>almost near</u>
alphabetical order	uniform distribution	100%	100%	50%	100%	17%	83%
random order	uniform distribution	100%	100%	75%	100%	67%	95%
	compound distribution	94%	~100%	~73%	~99%	—	—
	normal distribution	68%	95%	52%	84%	44%	75%

Figure(3). Transitivity Table For 1 Dimension

Suppose we are told A is near B, B is near C, and C is near D. This table shows how, no matter what case or method is used, our subjective estimates of the transitive nearnesses AC and AD correspond more to their almost near thresholds than to their near thresholds. Thus doubling the threshold to account for human suggestibility yields the correct result, which cannot be achieved by any kind of statistical sophistication

to a band of vagueness extending from about $.1 T$ to several times T (since the band ought to include almost all the probable location of the threshold). This means that the threshold varies by a large factor. This might be reasonable in situations where the meaning of near is ambiguous, rather than having merely its usual statistical vagueness.

The analogous values for the alphabetical case can be based on chi-square distributions which add almost as easily as normal distributions, changing only the number of degrees of freedom.

2 & 3 dimensions

For the two- and three-dimensional cases, we can use the multivariate case of one of the above distributions, to represent the probability that a given distance will be considered near. But the almost near threshold, being again twice the near threshold, gives an almost near area 3 or 7 times larger than the near area, unlike in one dimension, where they were equal. This corresponds to the fact that it is easier to pack objects closely in higher dimensional spaces. It also means that if an object has been judged near, the distribution of its location will no longer be the same shape (normalized) as the probability that a given location will be judged near (Figure 2a), which is independent of dimension, since Bayes' Law will give extra weight to the farther locations.

The probability that AC is near, using a bivariate uniform distribution, looks like roughly 67%, when AB is at its median near value of $T/\sqrt{2}$ (the radius which divides the near disc into two equal areas), as shown in Figure 4. This is slightly less than the random order case in 1 dimension (75%), and this seems reasonable.




Fig. (4) - Two-Dimensional Transitivity. The left large circle is the near threshold of A. If we are told B is near A, its median position will be on the dotted circle. If we are then told C is near B, it will be in the right large circle. The probability that C is near A is given roughly by the intersection of the two large circles.

USAGE

Martin's
work

Before an appropriate threshold can be computed, one must decide the sense in which near is being used in a sentence. In OWL <1975>, Martin points out that most common prepositions, such as "with," have up to a half dozen distinct, though related, meanings. Distinguishing the senses of a word forms a level of linguistic analysis intermediate to ordinary syntax and a full-fledged semantics which tries to explain in detail the meaning of each sense. If our threshold system is the full-fledged semantics for near, then this usage chapter is the intermediate level for near. I describe this level in terms of several usage dimensions.

symmetric
+
asymmetric

The first step is to decide whether usage is symmetric or asymmetric. In symmetric usage, the order of the two objects in the near phrase is reversible, and one is primarily interested in the size of the distance. In asymmetric usage, the first object is either more movable, more variable in location, or more uncertain as to location than the second object, which serves as a "base." The purpose is to provide some information about the first object, making use of the more well-understood base.

pair-definite

Usage is also either pair-definite or single-definite. In pair-definite usage, one is interested in the nearness of these two objects for some purpose that involves both of these specific objects. Neither object could be replaced by another; the identity of both objects is essential.

single-definite

In single-definite usage, the identity of the first object is essential, but the base is free to be chosen purely out of convenience for best conveying some information about the first object. It is selected by good general base qualities, such as nearness, familiarity, and lack of ambiguity.

semi-definite

Intermediate to pair-definite and single-definite usage are cases that may be termed semi-definite. Here, the second object's identity is essential only up to belonging in a certain set; as long as it is chosen from within this set, it is free to be selected by nearness or other general base criteria. Semi- and single-definite will be together referred to as indefinite cases. (I use the terms "case" and "usage" interchangeably.)

locative
+
comparative

In asymmetric case, we said the first object is the one whose identity is always essential, but the base is more well-known in location. Thus the base is providing some information about the essential object. This information is always one of two kinds. In the locative case, the listener knows the identity of the first object, and is being told its location. In comparative case, he knows the location of several candidates for first object, and is being told

DIMENSIONS OF LINGUISTIC USAGE OF NEAR, WITH EXAMPLES

23

		pair - definite	semi- definite	single - definite	
		specific identity of both objects essential	second object must merely be member of a certain class	second object free to be chosen by nearness of "good base" criteria alone	
SYMMETRIC		judges the <u>distance</u> between two objects	I can read easily, as my chair is near <u>the</u> lamp.	I can read easily, as my chair is near <u>a</u> lamp.	—
A S Y M M E T R I C U S A G E S	locative	<u>locates</u> first object	The chair is near the window, so it is drafty to sit in it.	Is the chair near the window or the door?	You'll find the red chair near the wall.
	comparative (selection)	<u>selects</u> first object from several candidates	He sat in the chair near <u>the</u> window, as he wanted the cool air.	He sat in a chair near <u>a</u> lamp, so he could read.	The chair near the wall is mine.
	comparative (comparison)	compares nearness of several first objects to the second object	The red chair is cooler than the blue, as it is nearer <u>the</u> window.	The red chair is cooler than the blue, as it is nearer <u>a</u> window	—

its identity, or else, several objects are being compared as to nearness to the base.

locative

If the locative case is single-definite, we are interested in the absolute location of the first object. If it is pair-definite, we are interested in the location of the first object with respect to the base. (The latter purpose is hard to distinguish from simply wanting to know the distance between the two objects, the usual purpose of symmetric usage, but locative is asymmetric because the first object is of less well-established position.)

comparative

If the comparative case is of the first kind defined above (identifying the essential object), it may be pair-definite or single-definite, but if the goal is comparing, it must be pair-definite, since there would be no point in comparing distances to an arbitrary object.

limitations

All these points will be made clearer by the following examples. They also suggest that our classification scheme is only a first approximation to what would be required to fully understand usage; for one thing, there are various kinds of semi-definiteness. More questions would have to be answered to determine how this usage information is to be used in understanding the meaning of the passage.

symmetric
examples

Examples of pair-definite symmetric usage are: "My home is near my job, so I don't waste much time travelling." "Salt and pepper are usually near each other on the table," or even "Put the salt near the pepper," if neither of them has been placed yet, so that the speaker is not specifying a location for the salt, but is merely saying "Wherever you put them, put them near each other."

A semi-definite symmetric example is "I can read easily, as my chair is near a lamp." This is symmetric because it can be reversed: "I can read easily because there is a lamp near my chair." The distance itself, rather than the location or identity of one of the objects, is our primary interest. If one wished to surmise which lamp the speaker was referring to, say in order to verify his claim, you would pick the one nearest the chair, or possibly a brighter and more noticeable one that was a little farther. Thus, even in symmetric usage, a non-essential object behaves something like a base.

True single-definite examples are hard to find in symmetric usage, since it is unlikely one would be interested symmetrically in distance to an arbitrary object, but we come a little closer in the example "Is there anything near the bomb?" ("Is the bomb near anything?") But this is still essentially semi-definite, since we are not interested in the base-like qualities of the arbitrary object, but are really referring to "anything valuable."

locative
examples

If the pepper were already on the table, and I wanted the salt

locative

examples

to be placed near it, or if I want the mustard placed near the cold cuts, this would be pair-definite locative. If I did not care what specific object the salt was near, but was merely using that object to specify location as in answering the question "Where is the salt? Is it near the cabbage?" (Reply: "No, it's near the spinach.") This would be single-definite locative. "Is it near the cabbage or near the spinach?" would be semi-definite case.

2-choice

is

semi-definite

The question of whether books on a shelf are near the left end of the shelf is implicitly semi-definite, since there is an implied comparison with the right end of the shelf. Thus, no matter how short the shelf is, a book closer to the left end will not be considered close to the right end. And in a long shelf, a book may be considered near the left end as long as it is clearly nearer the left end than the right end. (See "two-choice nearness," page 61, for a more detailed treatment of this problem.)

word

order

in

locative

case

As stated above for asymmetric cases, the first object is either more movable, more variable in location, or more uncertain as to location. Thus if a person is about to start up his car, I might say "Be careful. The cat is near the car." But if we are easing the car into the driveway, and we know the cat or Grandma is in the driveway, I might say, "Go slow. The car will be near the cat soon," or "The car is now near Grandma in her deckchair." However, if the driver didn't know the cat or Grandma were near the driveway, they would again have the more variable locations, and I might say "Go slow, the cat, or Grandma, is near the car."

inverts

in

word order

In buying a house, I would ask "Is the house far from the shopping center?" because, since I have not yet decided on this house, its position is somewhat variable, and even uncertain if I don't know where it is. But if I am with the realtor at the house I might ask "Is the shopping center far from here?" If I buy the house and my cousin comes to visit me for a week, he would ask "Is the shopping center far from your house?" since he knows my house better than he knows the shopping center. If he asked "Is your house far from the shopping center?" I might take this as an implied slur against my house's location, or as a hint that he did not want to do the shopping. This is because he is treating my house as variable, as something whose location is to be questioned, rather than as a given.

clues for
choosing
base

Further properties of a base are familiarity, importance, size and rarity. Rarity is useful because, with a common object, you don't know which one you're referring to (e.g. "His house is near the fire hydrant.") Familiarity and rarity are not mutually inconsistent, as you need only be familiar with this particular instance of the object.

comparative
examples

Consider the question "Is the red chair near the wall?" This sounds either pair-definite locative, expecting a reply like "No, it is far from the wall," or single-definite locative, expecting a reply

Comparative
examples

like, "No, it is near the door." However in some contexts one would be more likely to reply "No, the blue chair is near the wall." In this case, the task of the listener is not to locate a particular chair or learn its distance from the wall, but to select among, or compare, several chairs as to nearness to the wall, because one wants the chair(s) with this property. This is comparative usage. From the speaker's viewpoint, however, the chair was fixed (essential identity) and the wall as base was selected to describe it, something like in locative case.

Comparative
&
locative
grammar

Usage is more explicitly comparative when nearer than, nearest or as near as are used, or when near is in apposition to a noun, or is in an adjectival clause: "The chair near the wall is blue." or "The chair which is near the wall is blue." But if the clause is set off by commas, it is locative, being equivalent to a coordinate clause: "The chair, which is near the wall, is blue." (meaning: "The chair is blue, and it is near the wall.") Also locative are adverbial clauses: "You'll find the chair near the wall."

determining
usage:

What case is meant is not always obvious from the sentence, since a sentence may seem indefinite (e.g., "The chair is near the wall" - - the wall looks like it was chosen freely from among all room objects because it happened to be one of the nearest ones to the chair, and thus serves as a convenient descriptor of location), but the context may make clear why that base had to be chosen ("so you can lean backwards in it safely"), giving a symmetric pair-definite usage, or the context may specify a set from which the object had to be chosen (the previous sentence may have been "All the furniture in his room was either near the wall or near the window") giving a semi-definite locative.

definiteness

Objects referred to by "what?," "which?," "a," "any," etc., are clearly indefinite. If one object seems to be rather arbitrarily chosen, such as being a member of a set of identical or grouped objects, e.g. "My shot landed near the nineball," or if it is an uninteresting or immovable object, such as a wall or door (since their main use in nearness phrases is to locate or select other objects), then single-definite usage is likely.

symmetry

Symmetric usage is indicated when the two objects are similar. Symmetric and pair-definite locative are usually discriminable from other usages by the presence of a clear purpose in wanting to know the relative proximity of these two particular objects, e.g., "I hope the sprinkler isn't near the picnic table." Comparative case is usually explicit from the grammatical form of near or the relation of the near phrase to the sentence. However the case where the first object is a member of an identical or grouped set can be either locative or comparative, "Is the red chair near the wall?," as we saw at the beginning of the section on comparative usage. The cases depend on what question this sentence is construed to be in answer to. This

resemblance
to the
"reference"
problem

problem resembles the reference problem in language understanding.

after
determining
usage

Once usage has been diagnosed, the pair-definite cases are the simplest to deal with, since we know which two objects we are talking about, and neither can be upstaged by a ghostly member of some implicit class of objects. Nearness can be judged using a threshold determining system to be described later in this paper.

2-part
procedure for
indefinite
cases

But in the single-definite cases, one must first determine which, if any, of the eligible objects are near, and then select one, or more than one, if appropriate. The former procedure is basically the same as that to be used for the pair-definite determination of near, while the latter will be done as described presently.

strict
&
lax

thresholds

forced selection

A stricter interpretation of the threshold should be required for pair-definite cases, since, the two objects being already determined, our attention is likely to focus in sharply on them. In indefinite usage, on the other hand, we do not have to be quite as strict, since usually we are merely trying to find an object, and any hint that reduces the search is helpful. This is like the laxness, or "suggestibility" in transitivity on page 17. In semi-definite and comparative usages, it is often not even necessary for the object to be near, as long as it is nearer than the other objects, as in the end-of bookshelf example on page 25. This is especially true if the context absolutely requires selection of an object, e.g. "What city is nearest the North Pole?" But if a very near object exists, single-definite usage may actually be stricter, confining the choice to the nearest object to the exclusion of other near objects.

no near
object
or more
than one

Nevertheless it will usually be useful, even in comparative case, to apply the threshold. If I say "Take the nearest subway," you may reply "There is no subway near here." If I say "All the stores near here are closed," I certainly mean at least the nearest one, but it is unclear how many more I mean until some nearness criterion is applied. This may be based on grouping the stores by their relative nearness to each other, or may be based on a direct application of the threshold to each store. Clearly how we make our selection once we have applied the threshold will depend on how important it is to select an object, and what kind of results we are willing to put up with. (See "selecting the near object," below.) These matters require an analysis of the purpose of the judgment.

caution for
experimenters

A good example of pair-definite strictness occurred to me while exploring nearness thresholds empirically for this investigation. I found that when I asked myself "Is this object near that object?" I tended to use thresholds about half as large, or smaller, than thresholds that would pass in normal conversation, which is usually single-definite in referring to the undistinguished items I was using, such as books, pens and typewriters.

effect
of
vocabulary

Vocabulary is important in determining usage and strictness. "Near the end of the table" is not as near as "near the edge of the table." "End" emphasizes the range, being an implied semi-definite case, making a comparison with the middle. "Left end" would be a comparison with the right end. "Edge" emphasizes the object "edge" which has zero width, and hence is pair-definite with a very strict threshold.

selecting
the near
object

In answering the questions "Where is the squash?" or "Is the squash near the cabbage?" (probably single-definite locative) in a grocery store, one must not only check that they are near, but also whether something else might be nearer. In the latter case, and especially if the second object was only marginally near, or the third object much nearer, one must add a caveat: "Yes, it is near, but something else is nearer," or "Yes, in fact, it's right here near the eggplant." If second and third objects are both about the same distance, but the third object is more important, one favors it. If the third is equally important, one might mention both (but not too many, unless you can refer to them as a group). If the third is less important, ignore it. Note that in some ways immovable objects are better bases since you know where they are, but on the other hand this very fact makes them less remarkable and hence less likely to be used in some cases. Near a wall, door, or floorlamp possibly means nearer than near a pair of skis leaning against the wall, as the latter is more remarkable and so has a larger "sphere of influence."

more
suggestibility

If the second object was named, as in an assertive sentence, then one can "go along" with the choice in most of the above cases, but if it is we who are to select it, as in answering a question, then one follows the above rules strictly.

Thresholds for as near as, much nearer than, etc., will be developed as corollaries to our near determiner (see COMPARISON THRESHOLDS).

PART II: DETERMINING THE NEAR THRESHOLD

FACTORS INVOLVED IN DETERMINING THE THRESHOLD- - A MOTIVATION

prelude

Part II has a presentation consisting of several passes over the same subject matter. This overview chapter will introduce the reader to the complexity of the problem and the kinds of concepts that will be needed for the solution. The independent need for each of these concepts will be demonstrated. The following sections will develop each of these concepts in detail, and apply them to examples, to see how well they actually do. Often the same examples will be worked several times, at first using just one or two of the basic concepts, and later using the complete system, and the results will be compared.

definitions

Making the judgment of what is near in a given situation depends on a number of different kinds of factors. These include contextual information such as the purpose of the judgment, the range of possible values in the specific context (local range), the distance between other pairs of objects in the locality (local standard), and the size and shape of the objects involved; and general knowledge such as typical values of the range and standard in this kind of context, (global range and global standard), and the the scale of distances that characterise most human activities (absolute size). (In this essay, "context" will refer to information given in, or inferred from, the actual text or scene in which the particular instance of near being discussed appears, while "general knowledge" will refer to relevant general information culled from our long-term memory by the information in the context. The "situation" will refer to both. The relative importance of these factors itself depends on clues in the context, and on general knowledge about the context.

example

Consider the example: A friend of mine is in my room and asks me "Where is your copy of Brave New World?" I reply, "It's near the middle of the top shelf in my bookcase" or "It's near Blake, on my top shelf."

"information-
theoretic"
considerations

The purpose here is to help my friend find the book, given that he is in my room or knows how to get to it. Thus near should select out some reasonable fraction of the possible locations, i.e., the various positions on my top shelf. The selection should narrow the field of search by at least a factor of two or four, but not by so

"information-
theoretic"
considerations

much as to restrict the field to a region that could better be referred to by the terms "at" or "next to," as that would limit the applicability of the near concept, as well as make it a mere synonym for these other terms.

local range

Therefore an upper bound for the meaning of near would be some fraction of the length of the bookcase. The length of the bookcase is the local range. A lower bound would be some multiple of the size of the book. Not only does a smaller object need to have its location more exactly specified to be found easily, both because small objects are intrinsically hard to spot and because there can be more of them in the neighborhood, but also a small book can be merely near at a distance at which a large book would be "at" or "next to."

object size

linearity

The meaning of near should not only be confined somewhere within these bounds, but should depend continuously on both of them, if we are to make full use of the information they provide. We might expect that, given the ratio of these two bounds, the ratio of the near cutoff to each of them would be determined.

global
range

But this is not all. Even if we hold this ratio constant, by increasing the thickness of the books in proportion to the length of the shelf, we would be less likely to refer to a distance equal to, say 1/4 the length of the shelf, as near in a long bookshelf than in a short one. In a very short shelf, all the books would be near each other for some purposes (though they would not all be near the left end of the shelf). The effect of local range is moderated by its typicality or the global range. In other words, a theory which takes into account only local range and object size works for typical values of the local range, but, to the extent that the local range is atypical, its effect on the meaning of near is reduced. A typical range is one whose length is near the mean of the global range, the range of sizes of bookcase shelves as known from general knowledge. (Actually, I use the term "global range" to mean "typical range" thru most of this paper.)

absolute
size

But this is still not all. Suppose the local range is typical for the context. Our judgment will still be influenced by the absolute size of the distances involved. For a small context such as the inside of a change purse, we would be more inclined to refer to a given fraction of the range as near than we would for a large context such as a house, a city block, or a geographical area.

their
independence

But cannot the concept of absolute size also handle typicality of range? The atypically large range would automatically reduce the near fractional threshold, simply because it is large. This is not true, because a given size difference will have more effect if it is due to atypicality than if it is merely due to absolute size.

To see this, we will compare a naturally large distance with

room
crowdedness

one that is equal, but atypically large; the latter will seem larger because one tends to compare it with its typical value. Imagine a ten-foot square living room and a ten-foot square bathroom. One would say that the furniture in such a living room must be packed close together, whereas the furnishings in such a bathroom are far apart, even if these furnishings are the same in number and size.

standard

Finally, local and global standards may exist. On a library shelf from which three-quarters of the books have been removed (selected at random from the shelf), two books would be considered near, which, on a full shelf, would not be. Pairs of adjacent objects (objects which are each other's nearest neighbors along the line joining them) define a local standard for their locality, which is larger for a partly empty shelf than for a full shelf. As in the case of range, this local standard may be moderated by a typical, or global, standard, the typical distance for adjacent pairs in this kind of context. In the library example, it's not clear whether we should choose for global standard the distance between books on a typical partly empty shelf, or on a full shelf; probably the latter unless few shelves are full.

independence
of local
standard
and range
minimum

cars on a road

Note that a range may specify a minimum as well as a maximum, as in the cases of speed or distance between cars on a road. The range minimum is an absolute lower limit, while the standard is a typical value. What you mean when you say "the next car is following us rather closely" is determined by the commonest distance between successive cars within your sight (local standard), since you would be less likely to mention it if traffic was dense, and by the minimum distance cars can ever (or should ever) follow one another by on this kind of road (global range minimum). The latter increases the near threshold, as well as excluding small values from consideration.

discrete case:

small number
of objects

The above considerations must be modified when applied to situations where there are only a small number of objects. If my dresser has four or five drawers, I'm likely to refer to either of the top two as "near the top." The front two cars in a row of five are "near the front" of the line. Clearly objects twenty to forty percent of the way away from the end of a line would not be considered near that end if there were ten or more objects in the line. In the small number cases, we allow it because we do not want near to mean the same as "at."

large number
of objects

On the other hand, in the case of a large number of objects, such as a magazine of 75 pages, one may balk at calling many more than the last four or five near the end. Of course, information about number is not independent of our previous considerations; number can, in fact be deduced from range, object size, and standard distance. But this is a rather complicated calculation, since standard distance means the distance between the nearest points on two objects, not the distance between their centers. (This is necessary, as can be seen in

closely packed
or
evenly spaced

the examples of cars driving near each other, and two long pieces of furniture near each other.) It seems likely that number enters as an independent consideration, when the objects are closely packed or evenly spaced. This idea will form the basis of our treatment of the discrete case.

purpose

In the example of the bookcase, purpose did not play a major role. This was because the local range was explicitly given in the words "the top shelf of my bookcase." Were the range not specified, as in the cars on the highway, we would have to know the purpose of the judgement to help us to select a suitable range (in this case, a range minimum). Another example is: A pedestrian asks me "Is it far to Harvard Square?" Clearly I am more likely to answer "yes" than if a motorist asks me the same question.

summary

In sum, a procedure to compute near must involve both linear-like functions which weigh the effects of several continuously-varying measures, and heuristic rules with discrete thresholds or cutoffs, which tell which measures are relevant and what size they must be.

a
default
hierarchy

My suggestion is that the various factors be arranged in a flexible "default" hierarchy, or partial ordering, in which items initially at the same level, or in one order of precedence, may be put in another order of precedence if the situation advises it. Items at the same level get averaged in some continuous way, while items with precedence may either be given special weight in some continuous way (relative precedence), or may be given the whole say whenever their value is known or inferrable from the situation (absolute precedence), or when they, or other items, have certain values, or in certain qualitative situations (conditional precedence).

DETERMINING THE THRESHOLD - - CONTINUOUS CASE

Range and Standard

purpose

Once usage has been determined, one is ready to examine the factors that determine whether the two objects are indeed near. The first question to ask is "Is there a purpose explicit or inferrable from context? If not, from general knowledge?" Usually information from both sources will be needed. A purpose divides the distance domain into three regions: near, far, and in-between. A vaguer purpose may merely suggest that our threshold be "strict" or "lax."

range

Usually the context also suggests a range. Range and purpose are closely intertwined, each suggesting a suitable choice for the other. If both are present in the context, the range can be used to confirm the purpose. If one does not know the purpose, or if it is

unclear how to use it, he can usually divide the domain reasonably appropriately, from consideration of the range alone.

MBTA
example

If I see an ad "Apartment for rent, near MBTA," I infer the purpose "near enough to be easily accessible by MBTA." But this distance can be estimated simply by taking it to mean anything less than some moderately small fraction, say one-eighth, of the maximum distance an apartment in the metropolitan area is likely to be from the MBTA. If this seems too small, it is probably because the MBTA is designed so that most places are reasonably near it. If we really consider the whole metropolitan area in choosing the maximum, we will probably get a reasonable threshold. An alternative approach is to decide that the maximum for distance to the MBTA in the city is just not quite far, and interpolate the other CONSECUTIVE NEARNESS THRESHOLDS (page 67). This gives half the maximum as the near threshold (see also "standard", below, for another approach yielding this same result). See SCOPE AND METHOD for some further examples of range.

another
MBTA
example

If Mother tells her son, while waiting for the streetcar, "Do not stand near the tracks," the purpose can be inferred to be "near enough to get hit by the streetcar." But this distance can also be derived by taking it to be one-eighth of the platform width. In this example, however, the child would be more likely to assume Mother meant anything significantly less than the distance the other people were standing from the tracks, unless few other people were there. This would be an example of standard, which is often present in the context along with range. (See page 36 for more on standard.) In our apartment example, the distance of apartments from the MBTA is uniformly distributed through the entire range, so standard is not easily applicable, unless one defines it as half the typical maximum, giving the same answer as our modified range analysis in the previous paragraph.

standard

Let us define a threshold based on range as shown in Figure 6.




Fig. (6) - RANGE THRESHOLDS - The eighth of the range nearest 0 is near; the half farthest from it is far. Range refers to the range of possible distances between two objects in a given situation.

statistical
"justification"

Notice that if we consider the possible position of the object to be uniformly distributed within the range, as in our apartment near MBTA, the standard deviation is $1/12\sqrt{3} = .29$. Our threshold is .375 away from the mean, or about $4/3$ deviations. A threshold of $1/5$, such as Freiling uses, would be close to one standard deviation away from the mean.

2-dimensional
ranges

For a one-dimensional or elongated two-dimensional range, a threshold of $1/8$ the length of the range is adequate. But for a range with significant width, such as a square, this width should act to increase the threshold, since it increases the size of the range itself. However we wish to avoid complicated arithmetic such as computation of area or taking of square roots. One approach would be to define the effective one-dimensional range as the longest cross-section of the two-dimensional range (or a bit more), and use $1/8$ of this for the threshold. This is still impractical, since we usually know only the rectangular dimensions of rooms, not their diagonals.

square
ranges

A simple, adequate solution is, for squares, to define the effective one-dimensional range as twice S , the length of side, thus adjusting for the existence of the width. The threshold would then be $1/8$ of this effective range, or $S/4$. Figure 7 shows how nicely this works out for several examples. It is reasonable that the large square should have a threshold $\sqrt{2}$ times that of the small square, since their lengths of side are in that ratio. $\sqrt{2}$ also seems a reasonable ratio for the threshold of the small square to the threshold of the straight line.




Fig. (7) - TWO-DIMENSIONAL RANGES - IF

thresholds for near are always taken as $1/8$ of the range, then square ranges can be handled by defining their "effective 1-dimensional range" as twice the length of side.

rectangular
ranges

For rectangular ranges, we can usually get an adequate approximation by treating them either as a square or as a one-dimensional range, depending on how elongated they are. An intermediate model for 2×1 rectangles can be prepared also, if desired, by a linear interpolation of the two extreme cases.

example

Let us see what our rule does for an 8×1 rectangle and a square of equal area. The near threshold will be 1 for the rectangle, and $(2 \times 2 \sqrt{2})/8 = .7$ for the square, whose side is $2 \sqrt{2}$. The area near a point at the end of the rectangle is 1, but the area near a point at the edge of the square is $2 \times (.7)^2 = 1$ also, since this area extends .7 on both sides of the point, due to the large width of the square. Similarly, both figures have near areas of 2 for points near the center.

measuring
"shape"

We may see this equality as a balance of two opposing tendencies - - the square has smaller dimensions, reducing the threshold in proportion, but it is also more compact, leading us to expect that a greater portion of it will be near any point, other than a corner. Thus we are measuring shape without multiplying sides or taking square roots.

very long
ranges:

several
treatments

In the case of very long ranges, such as roads, or ranges in which one dimension is of different status than the other, an alternative definition may be indicated. Our regular definition treats long objects as one-dimensional, ignoring their width. This may be appropriate for comparing the location of buildings on one block of a street, where total length of the block is relevant, and width is not. But in talking about nearness of two objects on a sidewalk, or nearness of parked cars on a road, it is clear that the width of the range is more relevant as a near threshold than any function of the length, which is essentially infinite. The width is a natural scale unit for phenomena occurring in such a range. Sometimes one may even want to use a smaller scale, letting the width be range, and taking $1/8$ of it as threshold. This would be true if one has lost a small object on the sidewalk.

Examples where a single dimension is the range would be nearness of pages in a book, or of lines on a page.

Standard
threshold:

a) homogeneous
set of objects

If you have a collection of small bottles in a medicine cabinet, you may say "Put the aspirin near the antacid pills" meaning approximately the standard distance between adjacent bottles on the shelf. This is an adequate substitute for considering the actual purpose, which is that they be close enough so that when you see one, you see the other, but far enough to keep you from knocking one down when you pick up the other. If near is used in a sense that suggests "unusually near," as in "Why is that driver following us so closely?" then one should define near as significantly less than the standard distance between adjacent objects. "Significantly less" can mean

Standard and
"substandard"
nearness

"more than one standard deviation below the standard," but one-half the standard distance, which we used in our apartment-near-MBT example on page 34, is an adequate approximation.

b) distance
of many
objects of
one kind
to one of
another
kind

When you have a crowd of objects of one kind near an object of another kind, as people near a track in a subway station, you have a kind of standard which is a disguised comparative usage. There are again two cases. You may have to decide whether "near the track" means "where the nearest people to the track are," i.e., comparison with a standard analogous to the adjacency standard above, or "nearer than the nearest people," using the one-half adjacency criterion stated above. The latter case is often identifiable by use of words like "too near" or "so near."

RANGE EXAMPLES:

a) MIT campus

Let us try out our range and standard rules on some examples. If I tell someone on the MIT campus "Project MAC is near Polaroid," he will do well to assume I mean within one-eighth of the diameter of the campus (one mile), say 200 yards or less.

b) building

If we are on the eighth floor of the AI Lab building and I tell him Moore's office is near Kinsky's, he will do well to assume I mean about one-eighth of the greatest walking distance between two points on the eighth floor, i.e., at most 2 or 3 offices apart.

c) office

If we are discussing a 10 x 10 foot office, and I say "The desk is near the window," he could assume I meant "within about 2 1/2 feet of the window," 1/8 of the "effective one-dimensional range" for a square range (see page 35: dimensionality of range).

STANDARD
EXAMPLES

Let us now try the standard threshold on the above examples. The standard distance between adjacent buildings on campus is perhaps 100 feet or so, somewhat less than the range threshold. On the eighth floor, standard would mean one office away. But since offices are closely packed objects and we don't want near to mean "next to," (see page 31), we would expand the standard to two offices, the same as the range threshold. A more direct treatment of offices will be given under DISCRETE CASE. Inside the office, most objects on the order of size of a desk or window are, if "adjacent," about a foot apart probably, not much less than the range threshold of 2 1/2 feet.

Standard fairness can be defined as greater than four times standard nearness, as we shall do for all kinds of nearness.

distance
between
cities

(preliminary
treatment)

Let us try out the rules on distance between cities in the USA. Since the continental USA is about 3200 miles across, our far range threshold is 1600 miles, and our near range threshold is 400 miles. Thus Boston is near New York (215), and Detroit is near Chicago (262), but Boston is not near Washington (461), and L.A. is just outside the near threshold of San Francisco (412). Chicago is neither far nor near from New York (800). (Most figures taken from

distance
between
cities

Sunoco map of Eastern United States.) These judgments are a little on the near side, but probably fit well what a foreigner would say, and can be improved if we take into account the atypically large size, as a country, of the United States (global range), or the fact that our range may really be a part of the USA. Standards would give a slightly smaller threshold too, since the typical distance between adjacent large cities is closer to 200 miles (judging largeness by a local standard). If one is doing smaller cities, it would be likely that one's range is restricted to, say, New England, which again would give reasonable range and standard thresholds, as we shall see later, when we do cities with the completed theory (page 54).

books on a
bookshelf

Now let us try a smaller-sized domain, the bookcase. If it is 48 inches long, the range threshold is five inches, or about five or six books. Since, like with the offices, we have closely-packed objects, the standard must be increased to several books, so near won't be restricted to "next to." Thus we have rough agreement between range and adjusted standard rules. This example will be worked in detail under DISCRETE CASE.

lines on a page

A smaller domain yet is a page of a book. If a page has 48 lines, our range rule says the fifth line is near the top of the page. The adjusted standard would be close to that.

pages in a book

A still smaller domain is the distance between pages of a book. If a book has 400 pages, the range near threshold is 50 pages. Corrected standard gives several pages. Here we have significant disagreement. The implications of this will be discussed shortly.

selecting
range

Clues to selecting the appropriate range include what region the listener has already restricted the location to, the typical range for this kind of context, the distance of the speaker or listener from the scene, and any other distance playing a significant role in context.

selecting
standard

Standard should be based on other pairs of objects in the range, of the same type, or of the same or greater importance than the objects in question.

Object Size

limitations
of range
and
standard

Note that in all the cases where the range and standard rules worked well, namely the MBTA, MIT campus, bookcase, and lines-on-a-page examples, the near threshold was only a few times larger than the size of the objects involved. But in the pages-of-a-book example where the range rule overestimated, its threshold was 50 times the size of the objects involved (thickness of a page). Also, in the geographical example, where one could make a good argument that New York and Boston are not near, the near threshold (400 miles) was 20

times the diameter of a city (20 miles). This suggests that we can do better by taking into account object size.

object
size
threshold

Imagine two objects located in a void. Each has three dimensions. Consider any of these six dimensions to be large if it is greater than half the largest of the six. Figure 8 shows a table explaining the object size threshold we will use, and Figures 9a-n illustrate it pictorially, actual size, for small and medium size scenes.

absolute
size

The thresholds refer to the distance between the nearest points on the the two objects. This rule also takes into account absolute size, as you can see, taking the largest length of object: that a person can conveniently hold between two hands as the cutoff between medium and large. The largest size easily held between thumb and finger is the cutoff between medium and small. Of course one could add an interpolation rule to smooth the transition between categories. Without it, the error created will be as much as $\sqrt{2}$, but this is not much worse than our desired accuracy of 25% (see HOW VAGUE IS NEAR?). (It is a little worse than the $\sqrt{2}$ vagueness error though, since the object size probability distribution does not taper off at category edges like the normally distributed vagueness error does, but the error will be diluted when the object size threshold is meaned with the range threshold later.)

certain
omissions

The reader may note we have considered absolute size in our object size rule but not in our range rule. If the range rule is to be discriminating, its near and far thresholds clearly cannot vary too much from where I have set them. Also, we don't want to work too hard on absolute size. For a similar reason I have not used "global object size," though this omission may be unjustified. For further remarks on absolute size, see page 51. Also, see pages 51-52 for a way for absolute size to influence pure range judgments via "default" object size.

variety
of
treatment

Note that our treatment of dimensionality is an adaptation of our treatment of it in the range rule. There we doubled the effective length and threshold of a square because it has two large dimensions instead of one. Here we are doubling the threshold (on the average) whenever one more long dimension is added per object. The variation in this rule from smaller to larger objects is a way of moderating the influence of that sudden transition between size categories.

long objects

As in the case of long ranges, we may, for long objects (length more than eight times width, say; the width is near 0), wish to ignore the long dimension(s), and use the longest remaining dimension as the large dimension, if it is more suitable for our purpose. Or a geometric mean of several dimensions may be resorted to. Of course we want to limit use of these fudges to situations where we have a good excuse. One such excuse would be if not using

OBJECT SIZE THRESHOLDS

largest of all six dimensions < 1 inch			largest of all six dimensions is between 1 inch and $2\frac{1}{2}$ feet			largest of all six dimensions $> 2\frac{1}{2}$ feet		
total number of <u>large</u> dimensions in both objects	<u>near</u> thres- hold	<u>far</u> thres- hold	total number of <u>large</u> dimensions in both objects	<u>near</u> thres- hold	<u>far</u> thres- hold	total number of <u>large</u> dimensions in both objects	<u>near</u> thres- hold	<u>far</u> thres- hold
	(as fractions of the largest dimension)			(as fractions of the largest dimension)			(as fractions of the largest dimension)	
1	1	4	1	$\frac{1}{2}$	2	1	$\frac{1}{4}$	1
2,3,4,5	2	8	2,3,4	1	4	2,3	$\frac{1}{2}$	2
6	4	16	5,6	2	8	4,5,6	1	4

Figure (8) The Object Size Near Threshold is

obtained from the largest of the six dimensions possessed by the two objects whose nearness is being judged. We multiply this largest dimension by a power of 2 determined by the absolute size category and number of "large" dimensions of the two objects (those dimensions greater than half the largest).

then would put the near threshold near the edge of, or outside, the range.

poles

Let us illustrate by two examples. Near a 38-foot telephone pole is $7\frac{1}{2}$ feet by object size, and this is reasonable despite the great disparity between the pole's height and thickness, since the range is unlimited.

corridors

We will now examine a corridor problem, where we do have to drop a long dimension. We will show in addition, for this problem, that a satisfactory analysis is obtained only when a basic psychological characteristic of the situation is recognized, namely the indistinguishability of the two walls. For a corridor, both a range threshold and an object size threshold exist, and, anticipating the next chapter, we will use the spatial threshold, which is defined as their geometric mean.

we drop
the
long
dimension

Near a ten-foot high wall in a narrow corridor means within five feet by object size, since the wall has two large dimensions. We are ignoring the value of the (possibly) very long dimension parallel to the corridor, considering it only a second "large" dimension. If the width of the corridor is 4 feet, the range threshold is $\frac{1}{2}$ foot. The spatial threshold is then $\sqrt{5 \times \frac{1}{2}} = 1\frac{1}{2}$ feet. This is still too large, since it means an object only a half foot from the center of the corridor is near the wall. We could drop another long dimension, the height of the wall, but that does not seem a reasonable thing to do.

indistinguishability
of the two walls

An object is always within two feet of some wall, and for practical purposes the two walls are indistinguishable. When we say "He is standing close to the wall," we usually mean "close to one of the walls." This suggests taking two feet as our range. This yields a threshold of $\sqrt{5 \times \frac{1}{4}} = 1$ foot, better.

support
for this
idea

Let us check this by replacing one of the walls by a screen. Now it seems more reasonable that $1\frac{1}{2}$ feet is near the wall, or, on the other side of the corridor, $1\frac{1}{2}$ feet from the screen is near the screen.

The two wall case is really a question of "near a wall" versus "near the middle," like the "end of the table" example on page 28. The threshold is therefore stricter than in the screen case, where it is "near the wall" versus "near the screen," or simply "near the wall" versus "not near the wall."

another
approach

However the pure range threshold of $\frac{1}{2}$ foot may be at least as good a result here, suggesting that it is best to simply omit a threshold if its correct manner of use is not clear.

In each figure, the white objects are at the near threshold of the shaded object. The threshold, T , is given as a fraction of the largest dimension of the objects in question. (In our case, this will always be the largest dimension of the shaded object.) Diagrams are actual size. See Figure 8 on page 40 for explanation of threshold computation.

1 large dimension



a) medium absolute size (largest dimension > 1 inch). $T = \frac{1}{2}$




b) small absolute size (largest dimension < 1 inch). $T = 1$


Figures (9a-n)

OBJECT SIZE THRESHOLDS

2 large dimensions



c) medium absolute size, $T = 1$





d) small absolute size, $T = 2$

Figures (9a-n)

OBJECT SIZE THRESHOLDS

44

3 large dimensions
(split two-one among the two objects)



e) medium absolute size, $T=1$




f) small absolute size, $T=2$


Figures (9a-n)

OBJECT SIZE THRESHOLDS

3 large dimensions
(all in one object)



g) medium absolute size, $T=1$





h) small absolute size, $T=2$

Figures (9a-n)

OBJECT SIZE THRESHOLDS



46

4 large dimensions
(split two to each object)



(i) medium absolute size,

$$T=1$$





(j) small absolute size. $T=2$


Figures (9a-n)

OBJECT SIZE THRESHOLDS 47

4 large dimensions
(three in one object)



(k) medium absolute size, $T=1$





l) small absolute size, $T=2$

Figures (9a-n)



OBJECT SIZE THRESHOLDS

48

5 & 6 large dimensions



m) medium absolute size. $T = 2$



n) small absolute size.

for 5 dimensions, $T = 2$

for 6 dimensions, $T = 4$

Figures (9a-n)

Reconciling Different Thresholds

geometric
mean rules:

We are now faced with the problem of reconciling our several thresholds, as well as integrating local and global thresholds. If several thresholds are present and disagree, and the purpose does not suggest use of a particular threshold, then an appropriate threshold can be calculated as follows:

local-global
mean rule

1) If there is both a local and global range, take the geometric mean of their thresholds, and similarly for local and global standard.

spatial
threshold

2) If both range and object size thresholds can be formulated, and differ by more than a factor of about $1\frac{1}{4}$ to $1\frac{1}{2}$, take their geometric mean.

This new threshold may be called the spatial or indirect threshold, because both range and object size are based on spatial (geometrical) considerations as opposed to standard, which is a more direct kind of evidence based on comparison with analogous examples. Usually the spatial threshold will be a good substitute for the standard; if not:

two-mean
rule

3) If both a spatial threshold and a standard exist, and there is no good reason for choosing a particular one of them, take their GM.

A.M.
approximation

The arithmetic mean can be harmlessly substituted for the geometric when the thresholds differ by only a factor of 2 or so. Use of arithmetic mean for factors greater than 4 would tend to not reflect the order of magnitude of the smaller threshold.

relative
size of
thresholds

Note that either the range threshold or the object size threshold may be the larger. The former occurs when you have two small objects far apart, the latter when two large objects, or one large and one small object, are close together. In the latter cases, if they are very close (a range with small maximum), the spatial threshold will always yield a near judgment, since it is greater than the maximum distance. This occurs in judging whether two objects are close enough to rub each other. Here we must ignore object size. Similarly, when the objects are at great distances compared to their size, such as planets, the spatial threshold will likely yield far. If there are a great many objects, such as stars in the Galaxy, then neighboring objects are bound to be near by the range threshold, though they may still be far by object size threshold. (The spatial threshold here turns out to be about a tenth of a light-year, so that only members of multiple star systems are near.) Standard would be the only sensitive threshold in these cases. If one desires a discriminating judgment, one must consult the purpose for an

choosing a
discriminating
threshold

OUTLINE OF THRESHOLD DETERMINATION

50

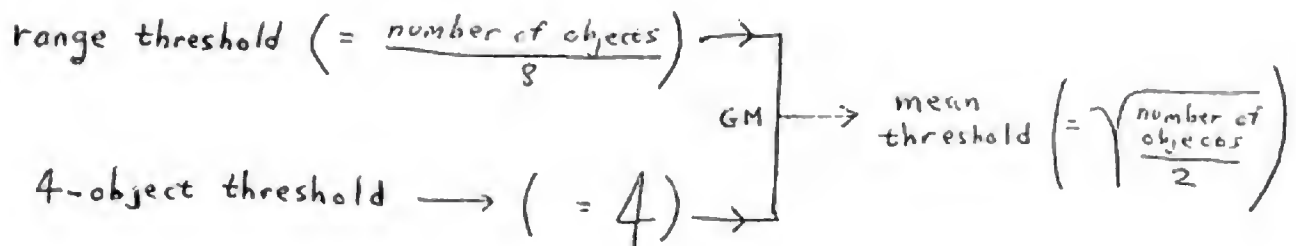
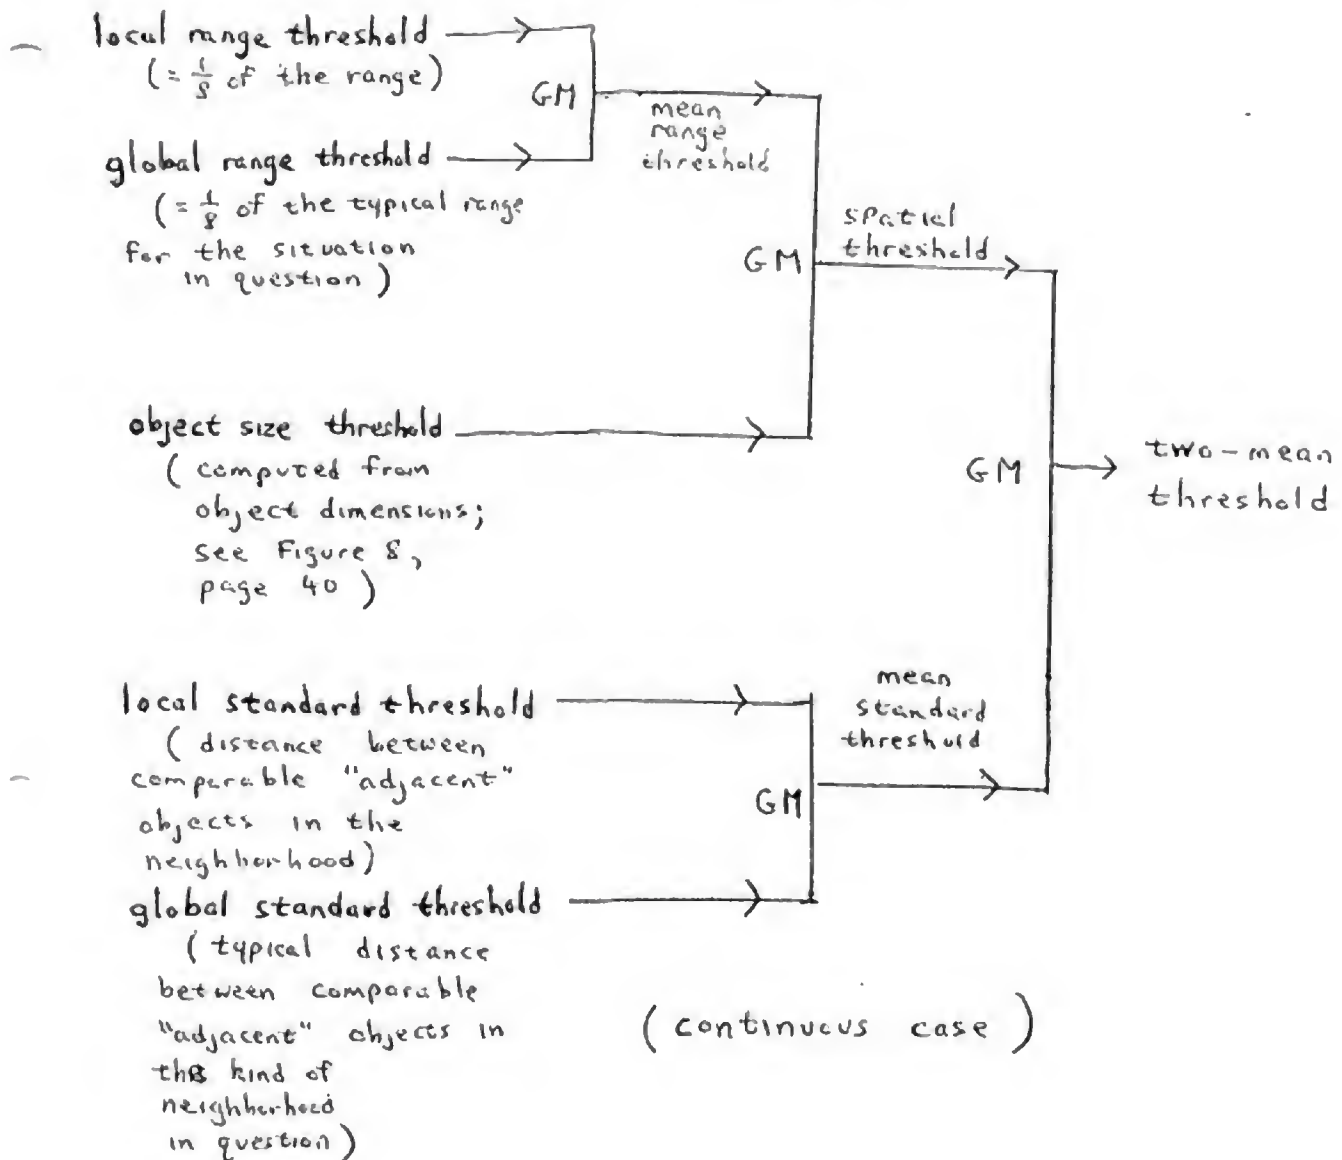


Figure (10)

(discrete case)

choosing a
discriminating
threshold

appropriate threshold, or failing that, use whichever rules discriminate. When range is much greater than object size, one is more likely to discard the object size threshold if one knows for sure that no objects fall within the spatial threshold, i.e., when range minimum exceeds it. This is the case for planetary orbits:

planetary
orbits

If one takes the radius of the solar system, three billion miles, as range, giving a near threshold of 400 million miles, and object size 10,000 miles (rough mean diameter of the planets), giving a near threshold of 10,000 miles, we get a spatial threshold of 2 million miles, so one might say of the solar system "No two planets are close, but the Earth and Moon are close." But if one already knows this, one might wish to use the pure range threshold of 400 million miles, and say "The inner planets are all close to each other, but not the outer planets." Standard is difficult to apply, since there are separate standards for the inner and outer planets. If one accepts the Titus-Bode Law as providing a standard, one can say "Mars and Jupiter are unusually far apart," since the ratio of their orbits is off Titus-Bode by a factor of 2.

default
thresholds

If there is any difficulty in the formulation or use of a particular threshold, it is usually best to simply omit it, as it is likely irrelevant. (E.g., the corridor problem on page 41.) If only a range rule or an object size rule can be formulated, we would like to use only that rule in calculating spatial threshold. This would be equivalent to assigning a default threshold of equal value to the other rule, corresponding, in the case of a pure range judgment, to a default size about equal to the range near threshold, and in the case of a pure object size judgment, to a default range several times the size of the two furthest points on the objects. The option to use default values gives modularity to our system, since we can do without missing information easily.

a slight
bug

However, there is a slight bug in this simple default idea. In the continuous case, using a pure range rule, the default object size threshold as defined above would be equal to the range threshold. Now, the ratio of the object size threshold to the object size varies from 1/4 to 4 over the nine different absolute size and shape categories. This means that we are implying different default object sizes, as a fraction of range, for different absolute size or shape categories. Large, globular objects are being assumed to have more restricted ranges than small, long ones. There is no reason why such an effect should occur in reality.

the
solution
has ...

The solution, fortunately, is simple and even has a good side-effect. One assumes, not a default object size threshold, but a default object size, equal to the range near threshold. For the intermediate absolute size category, with intermediate number of long dimensions, this yields an object size threshold equal to the range threshold. But for large objects, this yields a default object size

... a
good
side-effect

threshold $1/2$ as large as the range threshold, for intermediate number of long dimensions, and twice as large for small objects. We now compute the spatial threshold in the usual way, by taking the mean of the range and object size thresholds. The net result is to decrease large range thresholds by $1/\sqrt{2}$, and increase small ones by $\sqrt{2}$, thus providing a way in which absolute size can have some effect on pure range judgements, as we promised on page 39. This ratio is just large enough to be non-negligible.

why
mathematics
is
irrelevant
to AI

Our use of defaults means that the judgment of the distance between two geometric points will not be the limit one would get from a sequence of pairs of smaller objects, but rather corresponds to a pair of imaginary, reasonable-sized objects located at the points. It is clear that this discontinuity is justified, since points are usually used to represent the locations of reasonable-sized objects, not microscopic objects, and conformance to usage is what we want, not possession of nice mathematical properties such as continuity.

cheap
computation
of geometric
mean

It is important to note that our invocation of the GM rule does not mean we are committing ourselves to an expensive computation involving multiplication or square roots. Since we only want the results to ± 25 per cent, there ought to be a cheap method. To illustrate, we will show how an accuracy of $\pm 58\%$ can be achieved.

1 significant
binary digit

If we use only one significant binary digit of each number, rounding in the right direction, we will achieve this accuracy, since each rounded number is correct to within a factor of $3/2$. Squaring of this error factor in multiplication is undone when you take the square root. Since the product of two 1-digit numbers also is 1-digit, we lose no significant digits in multiplication. We must, however, use an accurate value of $\sqrt{2}$ (say 1.01 binary, good to about 13%, or 1.011, good to about 3%) when taking the square root, if the product has an odd number of zeros. Since no more arithmetic is done after that, these extra digits do not mean extra work. If we omit them, our result may be off by as much as $3/2 \sqrt{2} = 2.1$.

relative
precedence

more
planetary
orbits

The GM rule also gives us a handy way to compute relative precedence, if one ever needs it. In the astronomy example, for example, if we decide that range is more important than object size, but do not want to completely ignore object size, we can take the GM of the two as we did, and then take the GM of that value with the range threshold. This gives $\sqrt{2}$ million \times 400 million = 28 million miles. Thus adjacent inner planets are just about near each other, and Mars, but not Earth, is far from the sun. We get a similar result if we use the two-mean rule with a standard of 400 million, the mean distance between neighboring orbits. These results are relatively independent of our choice for object size (10,000 miles), since they vary only as the fourth root of it. We are, if you like, taking into account only the order of magnitude of the object size.

In examples like these, where two CM's have to be calculated in series, one would have to use more significant digits if the threshold is to be accurate to within the desired factor, but, on the other hand, since the the ratios are so large, we can afford a higher error factor in the threshold. In any case, this situation will be very rare.

absolute size,
qualitative
treatment

In our object size rule, we took some account of absolute size, dividing it into three categories. The most useful piece of advice for treating this subject, is to try to ascertain from the context whether the emphasis is on comparing the distances under discussion with normal-size distances, or with distances in their own size range (or some other size range). This is especially important for very large or very small distances. If the topic of discussion is just being introduced for the first time, or the listener is assumed ignorant, then one is likely to be talking absolute size. But once we are engrossed in the topic, and have more specific purposes in mind, a judgment relative to the standard for the topic is likely. Thus absolute size can be handled by selecting a range or standard that reflects the reader's expectation of what values are possible.

Nevertheless I will throw out several ideas for quantitative treatment of absolute size.

absolute size,
quantitative
treatment

For, especially, distances that do not differ too much from usual ranges, one can bias the threshold in a "strict" or "lax" direction by halving or doubling it. A more ambitious and less practical method, for astronomical and microscopic distances, would be to take a weighted geometric mean of the spatial threshold with a distance of 2 1/2 feet. The weight, or degree of relative precedence of the one factor over the other, could be selected somehow from one's estimate of their relative importance for the topic. Probably the spatial threshold should be given the major weight, since otherwise the threshold will not be very discriminating for the topic. Apart from arithmetic considerations, the main defect of this linear theory is its lack of refined consideration of any specific things we know about the purpose of the judgment, but it may be useful when we are ignorant and need something.

strict
& lax
thresholds

Earlier in this paper, we mentioned that a vague purpose, or a factor such as usage, may indicate that our threshold should be "strict" or "lax." This can be handled, as stated above, by doubling or halving the threshold. This is equivalent to moving the object to the next size category or number-of-large-dimensions category, in the case of object size threshold. Accordingly, the decision may be influenced by how close the object size, or the number of large dimensions, is to the edge of the category. This is a form of interpolation between category values.

Let us now see how the mean rules do for distances between

distance
between
cities:

USA

cities. The range near threshold was 400 miles. Since large cities are about 20×20 miles in extent (counting suburbs), except for New York which is 40×40 , a pair of large cities will have 4 "large" dimensions, and the object size near threshold is 20 miles. The spatial threshold is $\sqrt{400 \times 20} = 90$ miles. The standard distance is about 200 miles. Their mean is $\sqrt{90 \times 200} = 135$ miles (2-mean rule). This means Boston is not quite near New York (see CONSECUTIVE NEARNESS THRESHOLDS), Washington is just at the very near threshold of Baltimore (35), and Philadelphia is near New York (80) and Baltimore (85). Washington is not quite far from Boston, as is L.A. from San Francisco (just over 400 miles each). Detroit and Chicago (290) are near the not quite near-not quite far boundary. Cities more than 540 miles apart, such as New York and Detroit or Chicago are far, and cities over 1100 miles apart, such as New York and Denver, are quite far. Cities over 2200 miles apart are very far, such as New York and Phoenix, Arizona. Under a pure range rule, no cities would be very far, and only cities 3200 miles apart would be quite far. Doubling the range because of its two-dimensionality would only alter these results by a factor of $\sqrt{2} = 1.2$, or less, since the USA is really a 2×1 rectangle. I think our two-mean rule agrees pretty well with common usage, when the context is the whole United States. Usually, however, we use a regional context, because we are most interested in mileages when driving a car.

New
England

Let us look at the context of southern New England. Here the cities are about 10×10 miles in area, so the object size threshold is 10 miles. The range is about 200 miles, giving a range near threshold of 25 miles. The spatial threshold is $\sqrt{10 \times 25} = 16$ miles. The standard distance is about 40 miles, since "adjacent" pairs of cities, such as Boston-Providence, Boston-Worcester, Boston-Nashua, Providence-Worcester, and Hartford-New Haven all are 40 miles. (Springfield-Worcester is 50, but Springfield-Hartford is 26.) The two-mean rule gives therefore $\sqrt{16 \times 40} = 25$ miles for the near threshold and 100 miles for the far threshold. This happens to equal the range threshold. (This coincidence did not occur for large cities in the USA; corresponding values were: near two-mean threshold, 135; near range threshold, 400.) Thus Hartford-Springfield (26), Waterbury-New Haven (21), Bridgeport-New Haven (19), Providence-Fall River (18), and Boston-Brockton (18) are near. Boston-Springfield (89) and Springfield-Albany (91) are almost far. Boston-Portland (107) and Hartford-New York (113) are far.

New
York
City

Actually, since New York is 40×40 in area, we should expect it to have a bigger threshold. Its object size threshold for pairing with small New England cities is 20 rather than 10 miles. This gives it a two-mean rule near threshold of

$$\sqrt{\sqrt{20 \times 25} \times 40} = 30$$

or 30 miles for pairing with small New England cities, not much

distance
from
New York
City

bigger than our 25 mile threshold; due to the fact that the object size entered only as a fourth root. However we are saved if we remember that distance in CONTINUOUS CASE is measured from the nearest points of objects, not from their centers. Thus Norwalk (43) is near New York, since New York City has a radius of 20 miles; we can almost call Bridgeport near. Springfield (137 miles from center of New York) is near the far threshold. Thus most of the cities in western Connecticut, beyond Bridgeport, are neither near nor far from New York.

furniture

Let's apply our rules to furniture and walls in a living room. Two chairs 3 1/2 feet high and 2 feet wide will be near by object size if they are within 3 1/2 feet. If the room is 20 x 20 feet, the effective one-dimensional range is 40 feet and the range threshold is 5 feet. Ignoring the two-dimensionality of the range, we would get 2 1/2 feet. The standard distance also falls somewhere in this range, depending on whether one considers all chairs in the room, or all furniture in the room. Thus all three thresholds are similar and reasonable.

walls

A wall 10 feet high and 20 feet long is on the edge of the one-large-dimension and two-large-dimensions categories for object size. Let's interpolate object size threshold at 7 feet, the mean of 5 and 10 feet. Since the range threshold is 5 feet, we get a spatial threshold of 6 feet. These are large. The standard distance of furniture or people from a wall is probably about a foot, since the distance is likely 0 for furniture next to a wall but may be a few feet. But two feet from a wall is still near it. This figure can be obtained by taking the geometric mean of 6 and 1 according to the two-mean rule ($\sqrt{6} = 2.45$).

DETERMINING THE THRESHOLD - - DISCRETE CASE

Large Number of Objects

Before we do any more examples, we must deal with the "discrete" case of near, i.e., closely-packed or evenly-spaced objects, where the number of objects is more important than the distances or sizes.

a row
of objects

In the discrete case, we have a range, consisting of the total number of objects, or object positions, lined up in a row. Instead of object size or standard distance, we have as a measure of distance the number of object positions apart the two points in question are. Unlike distance between randomly spaced objects, discrete distance is measured from the centers of objects, not the nearest points.

4-object rule

Let us now consider several different measures of nearness. The strictest is the four-object rule. If two objects, or an object and a point, such as the middle or end of the line, are within four positions of each other, they are near. This rule is applicable to rows of a dozen or more objects. (With fewer objects, it may still be a good guess in the object-object case if the example is one where the range isn't clearly restricted to the row. E.g., if I ask "Where is Knuth?" and it happens to be at one end of a five book bookshelf at the other end of which is Human Problem Solving, one might answer "It is near Human Problem Solving." If the shelf is a little longer, we have a dilemma. Should we use the discrete rule which says "no," or the spatial threshold for two books in a room, which would still say "yes?" But let us restrict ourselves to the row and go on with the analysis of the discrete case.)

range rule

mean rule

The four-object rule is a rather strict rule in the case of very long rows, and lax for very short rows. A much laxer rule for large n is the pure range rule. Two entities are near if they are within $n/8$ positions of each other. For small n , this rule is too strict. We can compromise by taking the mean of the two rules. Round the result downwards. We summarize the behavior of these three methods, and a logarithmic rule as well in the table in Figure 11. Rules grow laxer from left to right. The log rule has been calibrated so as to equal the mean rule for $n = 10$. To test the rules, imagine a history paper of n pages, and someone asks you "Where does it mention Julius Caesar?" "Near the beginning" or "Near where it mentions Gaul."

n	4-object rule	mean rule	range rule	log rule
10	4	$[\sqrt{10}] = 3$	$1\frac{1}{4}$	$[2.2] = 2$
100	4	$[\sqrt{100}] = 10$	$12\frac{1}{2}$	$[4.4] = 4$
1000	4	$[\sqrt{1000}] = 32$	125	$[6.6] = 6$

Figure (11). Comparison of four rules for computing the near threshold in discrete case with n objects. The mean rule, which is the G.M. of the 4-object rule and range rule, gives the best results. (Numbers tell how many object positions apart two near objects can be.)

evaluation

From the table it is clear that the lpg rule is not the right one, and that the mean rule doesn't do badly, although for some purposes one might choose one of the other rules. Let's look at some examples that we considered earlier under the range and object size rules.

offices

in
a
building

Let us look at the offices on the eighth floor. There is a problem with the topology, since the offices are arranged along the periphery of a circuitous (rectangular) corridor. Let us consider the range to be the total number of offices in two adjacent sides of the rectangle, in analogy with the continuous treatment. This makes sense, since the distribution of distances from one office to the others is the same as if only half the rectangular corridor and offices existed, starting at that office. The range is about 28, the mean is $\lceil \sqrt{4 \times 28/8} \rceil = 3$. (Square brackets indicate the "largest integer in" function.) The continuous treatment gives $\lceil \sqrt{1 \times 28/8} \rceil + 1 = 2$. We added 1 since continuous distances are measured from the nearest points. Thus only the discrete rule makes Winston's office or MacDonald's office near Minsky's. This is perhaps reasonable, considering we are judging the whole eighth floor, not just the AI Lab here. If the continuous threshold of 2 offices seems better, perhaps this indicates that we are really judging doors, not offices, since the doors are all we can really see from the corridor. Doors, being not closely packed or (often) regularly spaced, would suggest continuous case. We will now look at some examples where the merit of the discrete rule is clearer.

pages

in
a
book

For nearness of pages in a 400-page book, we have a range threshold of $400/8 = 50$ pages. Let us see how well we can do with the continuous treatment first. We would get a threshold outside the range unless we ignore the size of the two large dimensions of each page. Considering them only as two more "large" dimensions along with the thickness, we have six large dimensions, small size category, and an object size threshold of one page thickness. The GM of 1 and 50 is 7, and adding 1, we get a spatial near threshold of 8, not altogether a bad value. But this approach looks artificial; it seems to make more sense to take a pure one-dimensional approach, since the page length and width are not only too large but also probably irrelevant in a more fundamental sense. In this case we have two large dimensions, the thicknesses of the two pages, and the object size threshold is 2, giving a spatial threshold of $\lceil \sqrt{2 \times 400/8} \rceil + 1 = 11$. The discrete rule gives $\lceil \sqrt{4 \times 400/8} \rceil = 14$. I think this is at least as good as the best of the continuous values. In any case, we see that the continuous values approach it more closely the more one-dimensional our approach is.

lines on
a page

Let us try nearness of lines on a page. Assume the page has 40 lines, each 60 characters long. The lines are 1 character wide, and the distance between adjacent lines is also about 1 character, so a line in the usual sense is 2 of our line widths wide. (For the

lines on
a page

purpose of the computation, we are taking a literal interpretation of the lines as objects, in the spirit of the continuous approach.) If we assume the characters have equal length and width, we get a threshold of $[\sqrt{1 \times 1/2 \times 60 \times 40/8}] + 1 = 13$, too large, like the corresponding method for pages in a book, which gave a threshold outside the range. Considering the lengths of the lines only as two more large dimensions together with the widths, we get $[\sqrt{2 \times 1/2 \times 40/8}] + 1 = 3$, much better. The pure 1-dimensional approach yields the same. The discrete rule gives $[\sqrt{4 \times 40/8}] = 4$.

books
on
a
bookshelf

Let us look at my bookcase, which is of typical length and contains 50 books on the top shelf, averaging 5/8 inch in thickness. The books are 8 inches or 9.6 book thicknesses high. The regular continuous rule gives $[\sqrt{(9.6) \times 50/8}] + 1 = 8$ books apart as near, once again rather large. Treating the length and width of the books only as two more large dimensions along with the thickness, we get $[\sqrt{1 \times 50/8}] + 1 = 3$, again much better, but perhaps a bit on the small side. The pure 1-dimensional treatment yields $[\sqrt{2 \times 50/8}] + 1 = 4$, the best value so far. The discrete rule gives $[\sqrt{4 \times 50/8}] = 5$.

evaluation

Although the appropriate variations of the continuous rule give reasonable results in all three of the above examples, in each case an equivalent result was provided much more directly and logically by the discrete rule. Let us look at two more bookcase examples to test out our local-global mean rule. We will compare the discrete and regular continuous rules.

long
bookshelves

Let's consider a bookcase 5 times as long. The discrete rule gives $\sqrt{4 \times 250/8} = 11.2$, the continuous $\sqrt{(9.6) \times 250/8} + 1 = 18.3$. But we forgot to consider global range. Since this bookcase is 5 times typical length, the global-local mean rule results in a decrease of threshold by a factor of $\sqrt[3]{5} = \sqrt{2.2} \approx 1.5$. Our discrete rule then gives $[11.2 / 1.5] = 7$, and our continuous $[18.3 / 1.5] + 1 = 12$ (since we must divide before adding 1). Again I think the discrete rule is better.

short
bookshelves

Now consider a short bookshelf with 10 books. The discrete and continuous rules give $\sqrt{4 \times 10/8} = 2.2$ and $\sqrt{(9.6) \times 10/8} + 1 = 4.5$. Corrected by global range, we get $[2.2 \times 1.5] = 3$ and $[3.5 \times 1.5] + 1 = 6$. The discrete rule seems just right.

suburbs
of a
city:

In judging the nearness of two suburbs of a city like Boston, the suburbs generally have to be abutting, or close to abutting, to be near. If there is a whole other suburb between them, they are only marginally near at most. E.g., Medford and Chelsea, separated by Everett; Medford and Cambridge, separated by Somerville; or Cambridge and Waltham, separated by Watertown. Watertown and Arlington are more likely to be considered near because Belmont, which separates them, is short from north to south, and perhaps because they are

demographically more similar than the other pairs.

by continuous

How do our rules do? Let's try the continuous treatment first. The range is the Boston area, diameter 16 miles. The range threshold is two miles. Object size is 3 miles, giving a threshold of 3 miles. (These distances are between the nearest points.) Their mean is just under 2 1/2 miles. This agrees well with our observations.

by discrete

But since suburbs are closely packed objects, we should use the discrete rule. 4 object lengths is $4 \times 3 = 12$ miles. The GM of 12 and 2 is 5. But since discrete distance applies to the centers of the objects rather than the nearest points, we must subtract twice the radius of a suburb, giving 2 miles as threshold, the same as the continuous value to the nearest mile. Since suburbs are so irregularly shaped, one can make as good a case that the continuous rule applies as that the discrete rule applies, so it is nice that they agree.

towns in
Mass.

Let's look at towns in Massachusetts. The range is 288 miles, giving 25 miles as range near threshold. Diameter of each is 4 miles. The continuous rule gives a threshold of $\sqrt{4 \times 25} = 10$. Discrete gives $\sqrt{(4 \times 4) \times 25} = 4 = 16$. This compares to 25 for cities in southern New England. One would expect cities to have the larger threshold, since they are farther apart.

Small Number of Objects

Let us see how the mean rule works for small n . Imagine a dresser or filing cabinet with drawers, a very short bookshelf, or a city block with several buildings.

near the
end and
middle of
short
rows

Since we measure from centers, adjacent objects are 1 apart, the end objects are 1/2 from the end, and in the even case, the central objects are 1/2 from the center, while in the odd case the central object is 0 from the center, etc. (see Figure 12).

Thus, for $n = 1$, the object is near both ends and the middle. For $n = 2$, each object is near its end and the middle, and marginally near the other object (reasonable, since they are neighboring, yet as far apart as possible within the range).

end strictness

For $n = 3$ and $n = 4$, only the end objects are near the end. This is rather strict, since we do not want near to mean "at." "The shirts are near the top of the dresser" could mean the top two drawers, even in a three-drawer dresser, certainly in a four-drawer dresser.

For $n = 5$ to about a dozen, the end object and their immediate

MEAN RULE, FOR SMALL N

n	mean rule threshold $\left(\sqrt{4 \cdot \frac{n}{8}} = \sqrt{n/2} \right)$
1	0.7
2	1.0
3	1.2
4	1.4
5	1.6
6	1.7
7	1.9
8	2.0
9	2.1
10	2.2

Figure (12). How the mean rule does for small n . For example, for $n=5$, the near threshold is $\sqrt{5/2} = 1.6$. Thus in a row of 5 objects, two adjacent objects are near, since their distance apart is 1.0, and no other two objects are near. But the second object is near the beginning of the row (distance 1.5). Distance is measured between centers of objects.

slightly
longer
rows

neighbors are near the end. Objects not adjacent can be near each other only at $n \geq 8$. This is possibly slightly too strict, but the threshold is close for $n = 7$. For all odd n , the objects neighboring the middle object are near the middle. Except possibly for $n = 3$, this is reasonable both in view of the "laxer than at" criterion, and because these objects are nearer the center than the ends. For $n = 3$ this is lax, as it would mean all objects were near the middle, which would seem to deprive near of any information content. Of course there would still be some information conveyed, since the probabilities of the end objects are less than that of the center object. For comparative and indefinite usages, these simple relative probabilities may be more useful than the cut and dried classifications we are making here (see page 27). For $n = 9$, the third and seventh objects are also near the middle, a little lax, though they are nearer to it than to the ends. For even n from 6 to about a dozen, the objects one removed from the middle pair are also near the middle. For $n = 6$, this is a little lax, but possibly acceptable; they are also, and equally as much, near the ends. For $n = 10$, we encounter for the first time objects which are near neither the middle nor the end.

middle
laxity

the reason
and a
patch

Note the three lax cases, 3, 6, 9; all occur in nearness to the middle. This is because the middle has objects on both sides of it, encompassing therefore, a lot of objects, but perhaps we feel that in order to convey at least one bit of information, not more than half the objects should be near the middle. We can amend the rule by dropping those objects outside the middle half. We can do the same in answering the question "Is it near (any) one of the ends?" similar to our choosing a stricter threshold in the problem of a man near the wall in a corridor (page 41). Such amendments, or caveats, to the general formula, based on common sense (possibly the desire to transmit a maximal amount of information) are very useful. We will soon see, indeed, how such commonsense ideas can form the basis of a definition of near.

prelude

The rest of this section, except for the last paragraph, is an attempt to analyze the relation between two-choice nearness (near one end vs. near the other) and three-choice nearness (near one end, near the other end, or near the middle). It tries to show why the $\sqrt{n/2}$ rule we have developed fits the 3-choice better than the 2-choice case, and by a "linear" argument, derives from this rule a new, possibly more fundamental, \sqrt{n} rule for the 2-choice case. I also develop an axiomatic approach for $n \leq 8$, based on commonsense properties of near. The 2-choice version agrees with the \sqrt{n} 2-choice rule, and the 3-choice version avoids the end-strictness and middle-laxity of the $\sqrt{n/2}$ 3-choice rule.

3-choice vs.
2-choice

If we are talking about nearness to one end versus the other, disregarding the middle, as in "Near which end of the street do you live?" we may wish to be laxer. For n less than about a dozen, an

a
2-choice
rule
by linear
approximation:

object is near one end unless it is almost at, or is beyond, the middle. In general, it is a good "linear" approximation to say, in the two-choice case, it is near one end if it would be near that end in the three-choice case for twice as many objects. I.e., the middle acts as an end for each half of the line. Actually, we need a little more strictness than in the 3-choice case, where an object can be near both the middle and an end for some n .

\sqrt{n}

Since $\sqrt{n/2}$ is the 3-choice formula, the 2-choice formula we get by this linear approximation is \sqrt{n} . The lowest n where the two- and three-choice rules disagree, apart from $n = 2$ and $n = 3$ where we thought the three-choice rule was a little strict, is $n = 7$. Results for both cases are shown in Figure 13.

why the
mean rule
($\sqrt{n/2}$)
fits the
3-choice
case
better
than the
2-choice
case

The reason why the mean rule fits the three-choice case better than the two-choice case is that it was derived from considerations that were pair-definite (or, looked at another way, single-definite), whereas the two-choice case is distinctly semi-definite. The usage vocabulary is a little confusing in this context, but let me explain. Near by the range or 4-object rules meant near in a pair-definite sense, because one is interested in nearness of a particular object to one particular point, the specified end. Thus we have a case of nearness to this point versus any other point, such as the middle or the end. This situation can also be described as single-definite (see pages 22 and 25 in USAGE), if we see ourselves picking out this end as a convenient base from among all points on the line, e.g., the middle and the other end. Since the range and four-object rules represent our general nearness theory, it is not surprising that the situation they fit can be viewed in several ways. The two-choice case, on the other hand, is distinctly semi-definite, as the near point must be chosen from a set of two specific points, the two ends.

\sqrt{n}
rule
more
fundamental

The \sqrt{n} rule fits the two-choice case possibly better than the $\sqrt{n/2}$ rule fits the three-choice case, despite the fact that we derived the former from the latter, using an approximation. This suggests the n rule is more fundamental, which makes sense, since both the two-choice case and the \sqrt{n} rule are simpler than the three-choice case and $\sqrt{n/2}$ rule, respectively. If we retain the \sqrt{n} rule, and if we were to replace the linear approximation by a more exact relationship, the 3-choice nearness-to-an-end rule would, according to the last observation in the paragraph on "3-choice vs. 2-choice," be slightly laxened, as we have wanted. This view is incorporated in the following axiomatic treatment.

axiomatic
treatment

The amendment idea of page 61 can be expanded into a complete axiomatic system for the small n cases. The axioms can be arranged in a precedence hierarchy, each taking effect only if it does not contradict the previous ones, or, equivalently for these axioms, if it is needed to disambiguate the meaning. I present the axioms for nearness to an end. Axioms for nearness to the middle in 3-choice

<u>n</u>	three-choice ($\sqrt{n/2}$)		two-choice (\sqrt{n})	
	threshold.	number of objects <u>near</u> the end	threshold	number of objects <u>near</u> the end
1	.7	1	1	1
2	1.0	1	1.4	1
3	1.2	1 (rather strict)	1.7	2
4	1.4	1 (rather strict)	2.0	2
5	1.6	2	2.2	2
6	1.7	2	2.4	2
7	1.9	2	2.6	3
8	2.0	2	2.8	3
9	2.1	2	3.0	3
10	2.2	2	3.2	3
11	2.3	2	3.3	3
12	2.4	2	3.46	3
13	2.5	3	3.6	4

Figure (13) COMPARISON OF 3-CHOICE AND 2-CHOICE RULES .

The number of objects near the end of a row of n objects depends on whether we mean "near the end vs. near the middle" (3-choice) or "near this end vs. near the other end" (2-choice). The latter case has a 1/2 threshold. Both cases are handled well by square root n - except for the two entries where the rule is a bit too strict.

case can be worked out similarly; the above-mentioned amendment enters at $n = 6$.

1) (information axiom) At least one, but not all, objects are near.

This handles the $n = 2$ case.

2) (plurality axiom) At least two objects are near (i.e. near does not mean "at").

This handles the $n = 3$ case.

3) (preference or semi-definiteness axiom) An object cannot be near one end if it is nearer the other end.

This handles the $n = 4$ case.

4) (excluded middle axiom) An object at the center, or adjacent to the center for even n , is near neither end.

This is adequate up to $n = 6$.

5) (2-choice/3-choice discrimination axiom) Any remaining undecided objects should be called near for 2-choice case, and not near for 3-choice case.

This is adequate up to $n = 8$.

These axioms make the 2- and 3-choice cases identical for $n \leq 6$, avoiding the strictness of the $\sqrt{n/2}$ rule for $n = 3$ and 4.

Of course there is a great deal of redundancy in these axioms; it seems almost pointless to have five axioms to explain only eight cases. Also, they seem to be phrased so as to be of as little use as possible for larger n . I found these drawbacks could not be avoided without either making the explanation of the small n cases unconvincing, or changing the results for the large n cases from what both the root rules and common usage prescribe. The axioms were not chosen to provide an efficient derivation of the meaning of near, but to show how its meaning is rooted in certain commonsense ideas. While information maximizing principles seem to be able to explain the meaning for very small n where there is little choice in the definition if it is to convey any information at all, they soon break down as they must, for near seems to be based rather on a concept of a "moderately small" distance, as we have seen in this paper. This concept is informationally efficient only in the context of a rich natural language vocabulary where it fills a small "ecological niche."

axioms
for
nearness
to an end
of a
short
row

evaluation
of
them

a
common
basis
for
discrete
&
continuous
rules

What common basis can we find for the discrete and continuous rules? Such a basis should show an analogy between the 4-object rule and the object size rule. The 4-object threshold is 4 times the distance between successive object positions, which is the smallest distance that plays a role in the discrete case. On the average, the object size threshold equals the object size, which is 4 times $1/4$ the object size, the smallest distance that plays a role in the continuous case. (It is the threshold for the large absolute size, one large dimension category.) Thus in both cases the object-dependent threshold is four times the smallest significant distance, or "quantum." This distance is at the very near object threshold, and hence can be described as a very small distance.

PART III: FURTHER DEVELOPMENTS AND APPLICATIONS

CONSECUTIVE NEARNESS THRESHOLDS

A hierarchy of thresholds (see Figure 14) may be defined in relation to the near threshold, which will be represented as 1. Ordinarily the ratio between consecutive thresholds should be 2. A ratio of 2 means that when we narrow a description down by one threshold we are adding one bit of information. According to the chapter HOW VAGUE IS NEAR? this is the closest the thresholds can be to each other without being mistakable for each other with any significant likelihood, and, if they were farther apart, there would be values in between which could not likely be taken as values of either threshold, thus suggesting that a new intermediate verbal judgment be invented.

a
ratio
of
2

rationale

These are default definitions in the sense that they apply when there is no good reason to think they don't apply. The purpose may suggest a stricter or laxer ratio than 2. In particular when the geometric mean threshold is several powers of 2 below the range threshold, we should increase the ratio, so as to spread the thresholds more evenly through the range. If we wanted to spread them exactly evenly, we should choose the ratio to be the fifth root of the ratio of range to near threshold, allowing us to interpolate the four higher thresholds between them. It is adequate to select convenient integer or fractional values for the thresholds, that yield ratios that are approximately equal. (See animal size example on page 78.)

when
to use
other
ratios

On the other hand, when the near threshold is simply the range near threshold, i.e., 1/8 of the range, only the range maximum is quite far, and nothing is very far. These latter characterizations are only applicable when the range is at least 16 or 32 times the size of the near threshold. See NON-MEASURABLE SPACES, page 72, for a supporting example.

extreme
categories
not always
applicable

When the range has a minimum greater than 8, a similar phenomenon may exclude the very near threshold. If we want to avoid this, as in cars travelling very near each other on a highway, meaning a carlength away (the minimum safe distance), we should compute all thresholds by their distance from the minimum, instead of their distance from 0.

CONSECUTIVE NEARNESS THRESHOLDS

<u>Name of Interval</u>	<u>Interval</u>
very far	$X \geq 16$
quite far	$X \geq 8$
far	$X \geq 4$
not quite far (almost far, nearly far)	$4 \geq X \geq 2$
not quite near (almost near, nearly near)	$2 \geq X \geq 1$
near	$1 \geq X$
quite near	$\frac{1}{2} \geq X$
very near	$\frac{1}{4} \geq X$
at	$\frac{1}{8} \geq X$

Figure (14). Consecutive Nearness Thresholds.

If the near threshold is represented as 1, then a sequence of other thresholds can be defined by multiplying it by positive and negative integral powers of 2. The terms in the left hand column then refer to intervals definable in terms of these thresholds. Notice that "not quite far" and "not quite near" refer to opposite sides of the same threshold, 2, the G.M. of near(1) and far(4). In some contexts, this threshold represents "standard" distance or nearness. The near threshold, one-half it, is then "substandard" nearness (pages 36-37).

$$\underline{at} = (\underline{near})^2$$

Note that if we think of near as a "small" distance, then our at threshold says that a second-order small distance is negligible.

parts &
sections

Sometimes we want a very lax threshold, as when we say "It's in the first part of the book." This is accomplished by using the word "part" which indicates a substantial portion of the book. "Part" could be taken as one-quarter to one-half the book. In this case the components of a part are not quite far from one another, by the range rule. A "section" could be one-eighth to one sixty-fourth of the book; its contents are near to one another, but not at the same place. This makes a section have at least several pages.

COMPARISON THRESHOLDS

In doing indefinite and comparative cases, we often have to compare the nearness of several objects, "comparands," to a reference object. If the purpose or standard does not specify thresholds, then we must fall back on a system based on range and object size, or range and number in the discrete case.

defined
in terms of
previous
concepts

We want to make the judgment depend as much as possible on concepts we have already developed. We can say that, basically, the two comparands are about, almost or nearly as near or as far as each other when the difference of their distances is within the near threshold of the two objects. They are as near or as far as each other when the difference is within their at threshold. The nearer one is much, or quite a bit, nearer than the other when its distance from the the reference object is less than that of the farther comparand by at least the ratio of consecutive nearness thresholds. See Figure 15.

choosing
range

Usually the two comparands have a common range which may be used in computing the near threshold. If the range is unknown, a good default might be the distance of the farther comparand, or a small multiple of it.

DIMENSIONS OTHER THAN DISTANCE

Kahn's
problem

Kahn in his master's thesis develops a time-specialist, but does not deal with concepts such as "recently" or "a while ago," which he points out are very context-dependent. These phrases mean essentially "near this time." Let us see how our near system can handle these concepts.




Figure (15) Comparison Thresholds. An object A, (not shown), is as far as, farther than, nearer than, etc. object B is from object C, when it is located at the corresponding distance on the scale showing fractions of the way from B to C. "As far as" extends $\frac{1}{8}$ on each side of 1.

Time:

a
"discrete"
example

Consider the statement "His birthday is near Christmas." The objects involved are days, which clearly match our discrete case. The range is $1/2 \times 365 = 183$, since no two dates can be more than half a year apart (similar to the offices in the circuitous corridor). The range threshold is $1/8 \times 183 = 23$ days. The 4-object rule gives 4 days. The GM of these is 9 or 10 days, a very reasonable answer.

a
"continuous"
example

Now consider "He came down with nephritis soon after his trip to Europe." This is continuous case, and the durations of the two events, the trip and the disease, resemble object sizes. If we take six weeks for the durations (it seems a good estimate for both of them), then we know the threshold must also be in this neighborhood, which is quite reasonable.

more
detailed
considerations

One may object that we should have halved the threshold because the durations are in what we could call the "large absolute size" category (more than a day) but this is balanced by the fact that there are two rather than one large dimensions (the smaller duration is more than half the larger). If the later event had been a shorter event like "He caught a cold," the threshold would be halved to three weeks because there is only one large dimension. A more detailed treatment might also take into account range (the upper limit on how much later he could have become ill, suggested, say, by the present date, or by some other event similar to the trip which would have been used as the reference if the illness had happened after it) and standard (a threshold for a cold might be shorter because one expects to catch them more often).

ANIMAL
SIZE:

by extremes

Let us see how our consecutive thresholds apply to the judgment of animal size. We will consider only macroscopic animals. Our range maximum is a whale, say 100 metres long. This gives a range near, or small, threshold of 12.5, say 10 metres. The range minimum, and, more important, the smallest distinguishable distance, is about 1 millimetre. If we equate this with the smallest significant distance (see page 65), we get an "object size" threshold of four millimetres. The geometric mean of that with the range threshold is about $2/10$ of a metre, or eight inches. Therefore a squirrel is a small animal. This is about as good as one can hope to get by this method, since the size of the largest animal is somewhat arbitrary.

by humans
as standard

A better idea is to take a human as a standard. To choose our ratio of consecutive nearness thresholds, the best idea is to decide what size would be, say, very large, and then interpolate to get large and quite large. But since we rarely see animals larger than the elephant, and are not very familiar with their sizes, it is best to base the interpolation on "elephant" as quite large leaving whales as "qualitatively" very large. (An alternative would be to decide what is the smallest we think of as large and extrapolate for quite large and very large.)

large
animals

An elephant is about 6 metres long. Call this quite large. Call humans (2 metres) standard, or not quite large (the GM of small and large). Then interpolate large as $3\frac{1}{2}$ metres, a rhinoceros perhaps. A horse is not quite large by this criterion, but it would be if we went by weight instead of length, since its shape is more compact than a human's.

small
animals

Let us interpolate small, quite small, and very small between humans (68") and our minimum visible size, say 0.1 inches. A factor of 5 will serve to separate these thresholds, giving small = 12" (cat), quite small = 2" (mouse), very small = 0.5 inches (insect). Note we have used a two-ratio system, a different ratio above and below the standard.

NON-MEASURABLE SPACES: EMOTIONAL CLOSENESS

an
extremely
simple
idea
works

In talking about how close I feel to various members of my class of 32 students, it is unlikely that I have a convincing numerical measure of distance. Nevertheless, if I assume that the closeness of students to me is uniformly distributed over a range extending from as close as possible to as far as possible, then I can say the closest four students (1/8 of the class) are close to me, the closest 2 of those are quite close, 1 is very close, the farthest 16 far, and the farthest 8 quite far. In a class of 32, no one is very far.

More likely, the distribution is chi-square with a number of degrees of freedom determined by the number of dimensions along which I can feel close to someone, e.g., intellectually, emotionally, etc.. But it is questionable whether such a sophisticated analysis is feasible or necessary. The point worth making is that the one-dimensional uniform distribution works pretty well.

GROUPING BY NEAR - - A FAILURE OF THE THEORY

an
attempt
to apply
the theory

Suppose we have n objects, located in some spatial configuration, and we wish to group them into groups based in some natural way on their nearness to one another. The approach suggested by our theory is to define two objects to be in the same group if they are near each other, using the diameter of the whole configuration as range. We then take the transitive closure of this nearness relationship, giving an equivalence relation which divides the objects into disjoint groups.

it fails

As the examples in Figure 16 show, this method is hopelessly inadequate, since in reality, an intra-group distance of as much as two-thirds the diameter of the whole configuration (Figure 16a, which is supposed to be almost a square) can exist, and this distance may be only slightly less than the inter-group distance. Our method doesn't even work in one dimension (Figure 16b). It seems clear that no method that could reasonably be considered an extension or modification of our theory could work.



Figure 16 - Our near theory can't handle even simple grouping problems, because significant differences in distance can be quite small.

a
slightly
better
idea

the true
complexity
of the
problem

An idea that works better is to sort the distances between "neighboring" objects into two classes, long and short, and then "join" objects separated by short distances, and take the transitive closure. This works for the above two examples, but would run into problems with more complex examples. One would have to decide when two objects are neighbors, and where the cutoff is between short and long. The latter is itself a one-dimensional grouping problem. Even if we could do this, the method would leave a lot to desire, since it does not take account of such important spatial considerations as where the distances being sorted are with respect to one another. What may be "long" in one part of the configuration may be "short" in another, and a group's shape is as important as its intra-group distance in deciding whether it is indeed one group. Clearly, a deep understanding of space is necessary to handle this problem.

Appendix: MATHEMATICAL CAPABILITIES OF THE HUMAN PERCEPTUAL MECHANISM

According to an article by S. S. Stevens in The Handbook of Perception, vol. II, <1974>, the well-known Weber-Fechner Law was replaced about 1953 by a more accurate power function law relating stimulus to response. The article leaves little doubt that perceptual mechanisms of the brain commonly compute square roots and other power functions. A particularly relevant example Stevens reports is that the subjective visual area of a square is proportional to the $2/3$ power of its actual area. The geometric mean of the area and the length of side is the $3/4$ power of the area. Stevens has the following to say:

" . . . Is it a general law, or does it hold only for vision and hearing? Experiments to answer that question have explored more than three dozen sensory and perceptual continua. The remarkable and quite unexpected result is that this psychophysical law seems to hold in all sense modalities. . . . In general, each sense modality has its own exponent, but the values of some of the exponents depend on such parameters as adaptation and contrast." - - (page 364)

" . . . The electrical recording of nervous activity turns out to give highly variable results, but in many sense modalities the electrical potentials have been shown to grow as a power function of the stimulus intensity. There appears to be little question, therefore, that the sensory systems are capable of power transformations. . . .

" . . . In some of the physiological experiments, the recorded data can be described by power functions approximately the same as the corresponding psychophysical functions." - - (page 368)

Acknowledgements

I would like to thank Gerry Sussman for suggesting the topic of this paper and encouraging me in pursuing this work. Michael Dunlavy provided some useful criticism, which led to my writing the chapter Some Remarks on Use of the Geometric Mean. The MIT AI Lab has very generously allowed me the free use of their facilities.

References

- Freiling, Michael; Frames and Functions, working paper, MIT, 1973
- Greenblatt, Richard D., Eastlake, Donald E. III, and Crocker, Stephen; The Greenblatt Chess Program, AI memo, MIT, 1969
- Hollerbach, John M.; Hierarchical Shape Descriptions of Objects by Selection and Modification of Prototypes, M.S. thesis, MIT, 1975
- Kahn, Kenneth M.; Mechanization of Temporal Knowledge, M.S. thesis, MIT, 1975
- Martin, William; DWL, unpublished, MIT, 1975
- Minsky, Marvin; A Framework for Representing Knowledge, AI memo, MIT, 1974
- Stevens, S. S.; Perceptual Magnitude and its Measurement, in The Handbook of Perception, vol. II: Psychophysical Judgment and Measurement, (Edward C. Carterette and Morton P. Friedman, editors), Academic Press, 1974
- Winston, Patrick H.; Learning Structural Descriptions from Examples, Ph.D. thesis, MIT, 1978